

SCHOOL OF MATHEMATICS AND STATISTICS,
UNIVERSITY OF SYDNEY

Stochastic Analysis Seminar on Filtering Theory

Author:
ANDREW PAPANICOLAOU
alpapani@maths.usyd.edu.au

These notes were originally written for the Stochastic Analysis Seminar in the Department of Operations Research and Financial Engineering at Princeton University, in February of 2011. The seminar was attended and supported by members of the Research Training Group, with the author being partially supported by NSF grant DMS-0739195.

Contents

1	Hidden Markov Models	4
1.1	Basic Nonlinear Filtering	5
2	Filtering and the VIX	9
2.1	Stochastic Volatility	9
2.2	The VIX	10
2.2.1	The VIX Formula	11
3	Stochastic Volatility Filter for Heston Model	14
3.1	The Filter	15
3.2	Extracting the Risk-Premium	18
3.3	Filtering Average Volatility in Fast Time-Scales	19
4	The Zakai Equation	21
4.1	Discrete Motivation from Bayesian Perspective	22
4.2	Derivation of the Zakai Equation	25
4.2.1	The Adjoint Zakai Equation	27
4.2.2	Kushner-Stratonovich Equation	28
4.2.3	Smoothing	28
5	The Innovations Approach	30
5.1	Innovations Brownian Motion	30
5.2	The Nonlinear Filter	32
5.2.1	Correlated Noise Filtering	34
5.2.2	The Kalman-Bucy Filter	35
6	Numerical Methods for Approximating Nonlinear Filters	37
6.1	Approximation Theorem for Nonlinear Filters	38
6.1.1	Weak Convergence	38
6.1.2	Consistency Theorem	40
6.2	Markov Chain Approximations	42

6.2.1	Approximation of Filters for Continuous-Time Markov Chains	42
6.2.2	Approximation of Filter for System of SDEs	44
6.2.3	Discrete-Time Observations	44
7	Linear Filtering	46
7.1	General Linear Filters	46
7.1.1	Reed/Matched Filters	47
7.1.2	Fourier Transforms and Bandwidth Filters	48
7.1.3	Wavelet Filters	51
7.2	Linear Gaussian Models	55
7.2.1	The Wiener Filter	55
7.2.2	The Kalman Filter	58
8	The Baum-Welch & Viterbi Algorithms	63
8.1	Equations for Filtering, Smoothing & Prediction	63
8.1.1	Filtering	64
8.1.2	Smoothing	65
8.1.3	Prediction	66
8.2	Baum-Welch Algorithm for Learning Parameters	67
8.2.1	Model Re-Estimation for Parametric Transition Probabilities	69
8.2.2	Model Re-Estimation for Finite-State Markov Chains	69
8.3	The Viterbi Algorithm	70
9	The Particle Filter	72
9.1	The Particle Filter	72
9.1.1	Sequential Importance Sampling (SIS)	73
9.1.2	Sampling Importance Resampling (SIR)	74
9.2	Examples	76
9.2.1	Particle Filter for Heston Model	77
9.2.2	Rao-Blackwellization	78
10	Stability, Lyapunov Exponents, and Ergodic Theory for Finite-State Fil-	
	ters	81
10.1	Main Ideas and Their History	81
10.1.1	The Counter Example	84
10.2	Stability for Markov Chain Models	86
10.2.1	Perron-Frobenius & Oseledec's Theorems	87
10.2.2	Lyapunov Exponents of the Filter	89

Chapter 1

Hidden Markov Models

We begin by introducing the concept of a Hidden Markov Model (HMM). Let $t \in [0, \infty)$ denote time, and consider a Markov process X_t which takes value in the state-space \mathcal{S} . We assume that the distribution function of X_t has either a mass or a density, and we denote mass/density with $p_t(x)$ such that

$$p_t(x) = \frac{d}{dx} \mathbb{P}(X_t \leq x) \quad \text{for densities,}$$

$$p_t(x) = \mathbb{P}(X_t = x) \quad \text{for masses}$$

for any $t \geq 0$ and $\forall x \in \mathcal{S}$. The generator of X_t is the operator Q with domain $\mathcal{B}(Q)$, such that for any bounded function $g(x) \in \mathcal{B}(Q)$ we have a backward equation,

$$\frac{\mathbb{E}[g(X_{t+\Delta t})|X_t = x] - g(x)}{\Delta t} \rightarrow Qg(x) \quad \text{as } \Delta t \searrow 0$$

for any $x \in \mathcal{S}$. Provided that regularity conditions are met, the adjoint leads to the forward equation,

$$\frac{d}{dt} p_t(x) = Q^* p_t(x).$$

Example 1.0.1. If $\mathcal{S} = \{x_1, \dots, x_m\}$ (a finite space) and the operator Q is a jump-intensity matrix such that $Q_{ji} \geq 0$ for all $i \neq j$ and $\sum_{i \neq j} Q_{ji} = -Q_{jj}$ for all $j \in \{1, \dots, m\}$. The forward equation is then

$$\frac{d}{dt} p_t(x_i) = \sum_{j=1}^m p_t(x_j) Q_{ji}.$$

Example 1.0.2. If $\mathcal{S} = \mathbb{R}$ and X_t is an Itô process such as

$$dX_t = a(X_t)dt + \sigma dB_t,$$

then $Q = \mathcal{L} = \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} \cdot + a(x) \frac{\partial}{\partial x} \cdot$ is the generator. The backward equation is then

$$\frac{d}{dt} \mathbb{E}g(X_t) = \mathbb{E}\mathcal{L}g(X_t) = \frac{\sigma^2}{2} \mathbb{E}g''(X_t) + \mathbb{E}a(X_t)g'(X_t)$$

for any bounded function $g \in C^2(\mathbb{R})$. Provided that $a(x)$ and the initial distribution satisfy some conditions for regularity, there is also a forward equation given by the adjoint

$$\frac{\partial}{\partial t} p_t(x) = \mathcal{L}^* p_t(x) = \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} p_t(x) - a(x) \frac{\partial}{\partial x} p_t(x) - a'(x) p_t(x)$$

for any $x \in \mathbb{R}$.

In addition to X_t , there is another process Y_t that is a noisy function of X_t . The process Y_t can be given by the SDE

$$dY_t = h(t, X_t)dt + \gamma(t, X_t)dW_t$$

where W_t is an independent Wiener process, or can be given discretely,

$$Y_{t_k} = h(t_k, X_{t_k}) + \gamma(W_{t_k} - W_{t_{k-1}})$$

where $(t_k)_k$ is a set of discrete times at which data is collected.

As a pair, (X_t, Y_t) are a Markov chain. The process X_t is of primary interest to us and is referred to as the ‘signal’ process, however it is not observable. Instead, the process Y_t is in some way observable and so we call it the ‘measurement.’ Hence, (X_t, Y_t) is an HMM and the goal is to calculate estimates of X_t that are optimal in a posterior sense given observations on Y_t .

1.1 Basic Nonlinear Filtering

Let \mathcal{F}_t^Y denote the filtration generated by the observations on Y up to time t . The optimal posterior estimate of X_t in terms of mean-square error (MSE) is

$$\hat{X}_t = \mathbb{E}[X_t | \mathcal{F}_t^Y] = \arg \min_{f \in \mathcal{F}_t^Y} \mathbb{E}(f - X_t)^2.$$

Proposition 1.1.1. \hat{X}_t is the unique \mathcal{F}_t^Y -measurable minimizer of MSE.

Proof. Let f be another \mathcal{F}_t^Y -measurable estimate of X_t . Then

$$\begin{aligned} MSE(f) &= \mathbb{E}(f - X_t)^2 = \mathbb{E}(f - \hat{X}_t + \hat{X}_t - X_t)^2 \\ &= \mathbb{E}(f - \hat{X}_t)^2 + 2\mathbb{E}(f - \hat{X}_t)(\hat{X}_t - X_t) + \mathbb{E}(\hat{X}_t - X_t)^2 \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}(f - \hat{X}_t)^2 + 2\mathbb{E} \left[(f - \hat{X}_t) \mathbb{E}[(\hat{X}_t - X_t) | \mathcal{F}_t^Y] \right] + \mathbb{E}(\hat{X}_t - X_t)^2 \\
&= \mathbb{E}(f - \hat{X}_t)^2 + \mathbb{E}(\hat{X}_t - X_t)^2 \\
&\geq \mathbb{E}(\hat{X}_t - X_t)^2 = \text{MSE}(\hat{X}_t)
\end{aligned}$$

with equality holding iff $f = \hat{X}_t$ almost everywhere. □

The filtering measure is defined as

$$\pi_t(\mathcal{A}) = \mathbb{P}(X_t \in \mathcal{A} | \mathcal{F}_t^Y)$$

for any Borel set \mathcal{A} , and for any measurable function g

$$\hat{g}_t = \mathbb{E}[g(X_t) | \mathcal{F}_t^Y].$$

Remark 1. *There are also smoothing and prediction distributions. When posteriors have density, we write*

$$\pi_{t|T}(dx) = \mathbb{P}(X_t \in dx | \mathcal{F}_T^Y).$$

*We say that $\pi_{t|T}$ is the **smoothing** density if $T > t$ and the **prediction** density if $T < t$. Smoothing requires significantly more calculation to compute, but the prediction simply requires us to solve the forward equation for X_t in the interval $[t, T]$ with initial condition π_t .*

Example 1.1.1. Filtering With Discrete Observations; The Bayesian Case. *Let \mathcal{S} be a countable state-space, let $h(x)$ be a known nonlinear function, let $\gamma(x) = \gamma > 0$, and for $k = 0, 1, 2, 3, \dots$ let there be specific times t_k at which observations are collected on Y_t . At each time t_k , let $Y_k = Y_{t_k}$ and denote the history of measurements up to time t_k as $Y_{0:k} = \{Y_0, Y_1, \dots, Y_k\}$. Denote $X_k = X_{t_k^-}$.*

Consider the following discrete differential:

$$Y_{k+1} = Y_k + h(X_{k+1})\Delta t_k + \gamma\Delta W_k$$

where $\Delta t_k = t_{k+1} - t_k$ and $\Delta W_k = W_{t_{k+1}} - W_{t_k}$.

Using Bayes rule, we find the the filtering distribution has a mass function

$$\pi_k(x) = \mathbb{P}(X_k = x | Y_{0:k})$$

for any $x \in \mathcal{S}$. It can be written recursively as follows,

$$\pi_{k+1}(x) = \frac{1}{c_{k+1}} \psi_{k+1}(x) e^{Q^* \Delta t_k} [\pi_k](x) \quad (1.1)$$

where $e^{Q^* \Delta t_k}[\cdot]$ is the kernel of X_t 's forward transition probabilities, $\psi_{k+1}(x)$ is the likelihood ratio of $\{X_{k+1} = x, Y_k\}$ given Y_{k+1} ,

$$\psi_{k+1}(x) = \exp \left\{ -\frac{h^2(x)\Delta t_k - 2(Y_{k+1} - Y_k)h(x)\Delta t_k}{2\gamma^2} \right\}$$

and $c_{k+1} = \int \psi_{k+1}(x) e^{Q^* \Delta t_k} \pi_k(dx)$ is a normalizing constant. Equation (1.1) can be shown to hold true through a use of Bayes formula and by the independence properties of the HMM.

Proof. (of equation (1.1)) Regarding the likelihood function of $\{X_{k+1} = x, Y_k\}$ given Y_{k+1} , it is

$$\begin{aligned} L(Y_{k+1}|Y_k, X_{k+1} = x) &\propto \exp \left\{ -\frac{1}{2} \left(\frac{Y_{k+1} - Y_k - h(x)\Delta t_k}{\gamma\Delta t_k} \right)^2 \right\} \\ &= \exp \left\{ -\frac{(Y_{k+1} - Y_k)^2 + h^2(x)\Delta t_k - 2h(x)(Y_{k+1} - Y_k)}{2\gamma^2} \right\} \end{aligned}$$

and since we are only interested in how this likelihood varies with x , we can remove the terms that do not have x in them,

$$\propto \exp \left\{ -\frac{h^2(x)\Delta t_k - 2h(x)(Y_{k+1} - Y_k)}{2\gamma^2} \right\} = \psi_{k+1}(x),$$

and so it is in fact the likelihood-ratio of $\{X_{k+1} = x\}$.

Now, by Bayes formula we have,

$$\begin{aligned} \pi_{k+1}(x) &= \frac{\mathbb{P}(X_{k+1} = x; Y_{0:k+1})}{\mathbb{P}(Y_{0:k+1})} \\ &= \frac{\mathbb{P}(Y_{k+1}|X_{k+1} = x; Y_{0:k})\mathbb{P}(X_{k+1} = x; Y_{0:k})}{\mathbb{P}(Y_{0:k+1})} = \frac{\psi_{k+1}(x)\mathbb{P}(X_{k+1} = x; Y_{0:k})}{\mathbb{P}(Y_{0:k+1})} \\ &= \frac{\psi_{k+1}(x) \sum_{v \in \mathcal{S}} \mathbb{P}(X_{k+1} = x, X_k = v; Y_{0:k})}{\mathbb{P}(Y_{0:k+1})} \\ &= \frac{\psi_{k+1}(x) \sum_{v \in \mathcal{S}} \mathbb{P}(X_{k+1} = x|X_k = v)\mathbb{P}(X_k = v; Y_{0:k})}{\mathbb{P}(Y_{0:k+1})} \\ &= \frac{\psi_{k+1}(x) \sum_{v \in \mathcal{S}} \mathbb{P}(X_{k+1} = x|X_k = v)\mathbb{P}(X_k \in dv|Y_{0:k})}{\mathbb{P}(Y_{k+1}|Y_{0:k})} \\ &= \frac{\psi_{k+1}(x) \sum_{v \in \mathcal{S}} \mathbb{P}(X_{k+1} = x|X_k = v)\pi_k(dv)}{\mathbb{P}(Y_{k+1}|Y_{0:k})} \end{aligned}$$

$$= \frac{\psi_{k+1}(x) e^{Q^* \Delta t_k} \pi_k(dv)}{\mathbb{P}(Y_{k+1}|Y_{0:k})}$$

and clearly, $\mathbb{P}(Y_{k+1}|Y_{0:k})$ is the integral of numerator of the last line over x . \square

The Bayesian filter is an essential tool for numerical computations of nonlinear filtering. The Kalman filter (see Jazwinski [28]) is also an important tool, but it only applies to linear Gaussian models or models that are well-approximated as such. The contemporary way to compute nonlinear filters is via Monte Carlo with a particle filter (which we'll talk about in a later section). In continuous time, approximating filters based on discretization of the differential dY_t have been shown to converge as $\Delta t \searrow 0$ for a certain class of filtering problems, but we must be able to approximate the law of X_t , and h must be bounded (see Kushner [33]).

In summary, an HMM consists of a pair of process (X_t, Y_t) where X_t is an unobserved signal which is a Markov process, while Y_t is an observable process that depends on X_t through a system of known functions and known parameters. We use filtering to compute the posterior distribution of X_t given \mathcal{F}_t^Y .

Chapter 2

Filtering and the VIX

2.1 Stochastic Volatility

Consider an equity model with stochastic volatility,

$$dS_t = \mu S_t dt + f(X_t) S_t dW_t$$

where S_t is the price of a stock, the function $f(x)$ is known and X_t is a hidden Markov process. For instance, the Heston model, where $f(x) = \sqrt{x}$ and

$$dX_t = \kappa(m - X_t)dt + \gamma\sqrt{X_t}dB_t$$

where we model the volatility leverage effect by saying that $\frac{1}{t}\mathbb{E}B_tW_t = \rho$ with $\rho \in [-1, 0)$.

In general, if we observe a continuum of prices, then $f(X_t)$ is measurable with respect to the filtration generated by $\{S_\tau : \tau \leq t\}$. Let $Y_t = \log S_t$, and notice that

$$dY_t = \left(\mu - \frac{1}{2}f^2(X_t) \right) dt + f(X_t)dW_t.$$

For a fixed $t > 0$, let $(t_k)_k$ be a partition of $[0, t]$, then the quadratic variation of Y is the cumulative variance

$$[Y]_t = \lim_{\|P\| \searrow 0} \sum_k (\Delta Y_{t_k})^2 = \int_0^t f^2(X_\tau) d\tau \quad \text{in probability}$$

where $\|P\| = \sup_k (t_{k+1} - t_k)$. Clearly, then $\int_0^t f^2(X_\tau) d\tau$ is \mathcal{F}_t^Y -measurable, and if $f(X_t)$ is a continuous process we have

$$\frac{d}{dt}[Y]_t = f^2(X_t)$$

is also \mathcal{F}_t^Y -measurable. So at the very least (e.g. for $f(X_t)$ a continuous process) volatility is observable for almost everywhere t , and X_t is observable if f^{-1} exists.

Nonetheless, it is still beneficial to have a Markov structure for X_t so that we can price derivatives on S_t . For instance, in the example by Elliot [22], the Black-Scholes price of a European call option in the presence of Markovian volatility, when $\rho = 0$ is

$$C(t, S_t, X_t; T, K) = \mathbb{E}^*[C_{BS}(t, S_t; T, K, Z_{[t,T]})|X_t, S_t] = \mathbb{E}^*[C_{BS}(t, S_t; T, K, Z_{[t,T]})|\mathcal{F}_t^Y]$$

where $Z_{[t,T]} = \frac{1}{T-t} \int_t^T f^2(X_s)ds$, and $\mathbb{E}^*[\cdot]$ is the market's pricing measure. Given X_t and the parameters of X 's dynamics under the market measure, we can compute the expected return of the call option either explicitly or through Monte Carlo.

If X_t is not known (i.e. observations are discrete) then we could take a filtering expectation,

$$C(t, S_t; T, K) \stackrel{?}{=} \int C(t, S_t; T, K, x) \pi_t(dx)$$

but such a price would be significantly biased if the market placed any premium on volatility. This option pricing formula exemplifies the challenge of interpreting the filter in financial math.

2.2 The VIX

Filtering can be used to extract the risk-premium placed on volatility by the market. Let S_t be the S&P500 index. For any time t and some time-window $T > 0$, the VIX index is the square root of the market's prediction of average variance during $[t, t+T]$,

$$V_t = \mathbb{E}^* \left[\frac{1}{T} \int_t^{t+T} f^2(X_s)ds \middle| \mathcal{F}_t^m \right]$$

where $\mathbb{E}^*[\cdot]$ is the market's pricing measure, and \mathcal{F}_t^m is the filtration generated by *all* the information in the market (i.e. $\mathcal{F}_t^Y \subset \mathcal{F}_t^m$). The process V_t is the 'fair' price of a variance swap whose floating leg is the realized variance,

$$RV_{[t,t+T]} = \lim_{\|P\| \searrow 0} \frac{1}{T} \sum_k (\Delta Y_{t_k})^2 =_p \frac{1}{T} \int_t^{t+T} f^2(X_s)ds$$

where $(t_k)_k$ is a partition of $[t, t+T]$ (i.e. $t = t_0 < t_1 < \dots t_N = t+T$, with $\|P\| = \sup_k (t_{k+1} - t_k)$ going to zero as N gets large). For a pre-specified notional amount, the payoff of a variance swap is

$$\text{notional} \times (RV_{[t,t+T]} - V_t).$$

2.2.1 The VIX Formula

For diffusion models without jumps, it was shown by Demeterfi et al [19] that a portfolio of out-of-the-money call and put contracts along with a short position in a futures contract replicates the VIX. It was shown by Carr et al [14, 15] that models with jumps can be approximated by the same setup. The following lemma derives the strategy for diffusions:

Lemma 2.2.1. *Let T denote the life of the contract, let $F_{t,T}$ denote the future price on S_{t+T} at time $t \geq 0$, and let r be the rate so that $F_{t,T} = S_t e^{rT}$. If S_t is purely a diffusion process (i.e. has no jump terms in its differential), then the market's expectation of future realized variance is*

$$V_t = \frac{2e^{rT}}{T} \left(\int_{K \leq F_{t,T}} P_t(K, T) \frac{dK}{K^2} + \int_{K \geq F_{t,T}} C_t(K, T) \frac{dK}{K^2} \right) \quad (2.1)$$

where $P_t(K, T)$ and $C_t(K, T)$ denote the price of a put and a call option at time t with strike K and time to maturity T , respectively.

Proof. Under the market measure, there is a Wiener process dW_t^* such that the returns on the stock satisfy

$$\frac{dS_t}{S_t} = rdt + f(X_t)dW_t^*$$

and the log-price satisfies

$$d \log(S_t) = dY_t = \left(r - \frac{1}{2} f^2(X_t) \right) dt + f(X_t) dW_t^*.$$

Integrating the returns and the log-price separately and then subtracting, we can eliminate all randomness to get the cumulative variance,

$$\int_t^{t+T} \frac{dS_\tau}{S_\tau} - \int_t^{t+T} dY_\tau = \frac{1}{2} \int_t^{t+T} f^2(X_\tau) d\tau$$

which shows us that the realized variance satisfies the following:

$$\begin{aligned} RV_{[t, t+T]} &= \frac{2}{T} \left(\int_t^{t+T} \frac{dS_\tau}{S_\tau} - \log(S_{t+T}/S_t) \right) \\ &= \frac{2}{T} \left(\int_t^{t+T} \frac{dS_\tau}{S_\tau} - \log(S_{t+T}/F_{t,T}) - \log(F_{t,T}/S_t) \right). \end{aligned} \quad (*)$$

Then, through some simple calculus, we see that

$$- \log(S_{t+T}/F_{t,T})$$

$$= -\frac{S_{t+T} - F_{t,T}}{F_{t,T}} + \int_{K \leq F_{t,T}} (K - S_{t+T})^+ \frac{dK}{K^2} + \int_{K \geq F_{t,T}} (S_{t+T} - K)^+ \frac{dK}{K^2}.$$

Plugging this expression for $-\log(S_{t+T}/F_{t,T})$ into (*) we see that the realized variance can be written as the payoffs of several contracts and a continuum of puts and calls that were out-of-the money at time t :

$$\begin{aligned} RV_{[t,t+T]} &= \frac{2}{T} \left(\int_t^{t+T} \frac{dS_\tau}{S_\tau} - \frac{S_{t+T} - F_{t,T}}{F_{t,T}} - \log(F_{t,T}/S_t) \right. \\ &\quad \left. + \int_{K \leq F_{t,T}} (K - S_{t+T})^+ \frac{dK}{K^2} + \int_{K \geq F_{t,T}} (S_{t+T} - K)^+ \frac{dK}{K^2} \right). \end{aligned}$$

Taking expectation of both with respect to the market measure, the noise in $\int \frac{dS}{S}$ vanishes, and since both $\log(F_{t,T}/S_t) = rT$ and $\mathbb{E}^* S_{t+T} = F_{t,T}$, we have

$$V_t = \mathbb{E}^*[RV_{t,T} | \mathcal{F}_t^m]$$

$$= \frac{2}{T} \left(\int_{K \leq F_{t,T}} \mathbb{E}^*[(K - S_{t+T})^+ | \mathcal{F}_t^m] \frac{dK}{K^2} + \int_{K \geq F_{t,T}} \mathbb{E}^*[(S_{t+T} - K)^+ | \mathcal{F}_t^m] \frac{dK}{K^2} \right)$$

and if we multiply and divide the RHS by e^{-rT} we get the result. \square

Risk Premium

Under the market measure, the expected returns on a variance swap are zero. However, statistically speaking, variance swaps exhibit a slight bias against the holder of the contract. In other words, the person who receives $notional \times (RV_{[t,t+T]} - V_t)$ at time $t + T$ will have an average return that is slightly negative. But because of the volatility leverage effect, the variance swap has the potential to provide relief in the form of a positive cash flow when volatility is high and equities are losing. When quoted in the market, the VIX is quoted as the square-root of V_t in percentage points, and is a mean-reverting process that drifts between 15% and 50%. The VIX has the nickname ‘the investor fear gauge’, as it should because it is composed primarily of out-of-the-money options which means that there is an increase in crash-a-phobia whenever the VIX increases.

To get a more precise understanding of investors’ fears, it would be nice to remove any actual increases in volatility and merely examine the bias in the market measure’s prediction of variance. In other-words, we would like to predict variance in the physical measure, and then compare it with the VIX’s prediction to gain a sense of how much of a premium is being placed on risk, or how much fear there is out there.

If we can identify an HMM whose observable component generates \mathcal{F}_t^m , then the market's price of volatility risk (aka the risk-premium) is

$$\mathcal{RP}_t = V_t - \mathbb{E} \left[\frac{1}{T} \int_t^{t+T} f^2(X_s) ds \middle| \mathcal{F}_t^m \right].$$

Remark 2. Carr and Wu [16] point out that if we define the martingale change of measure with the Radon-Nykodym derivative M_t , then

$$V_t = \mathbb{E}_t^*[RV_{t,t+T}] = \frac{\mathbb{E}_t[M_{t+T}RV_{t,t+T}]}{\mathbb{E}_t[M_{t+T}]} = \mathbb{E}_t[RV_{t,t+T}] + cov_t \left(\frac{M_{t+T}}{\mathbb{E}_t[M_{t+T}]}, RV_{t,t+T} \right)$$

where $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_t^m]$ and $\mathbb{E}_t^*[\cdot] = \mathbb{E}^*[\cdot | \mathcal{F}_t^m]$, which leads to an expression for the risk-premium, and is \mathcal{F}_t^m -measurable provided that M_{t+T} is known.

More on the VIX

Historically, indices like the S&P500 exhibit contrary motion with volatility. In particular, periods of high volatility often coincide with bearish markets. Whaley [46] describes how tradable volatility assets, such as the VIX, provided market makers with new ways to hedge the options they had written. For instance, a market maker who was short a portfolio of options has always been able to go long in some other types of contracts to reduce the portfolio's Delta to almost zero, meaning that the portfolio would not be hugely affected by small changes in the value of the underlying. When VIX was introduced, it allowed the same market maker the opportunity to also hedge the Vega of his/her short position in options. Prior to volatility hedging instruments, a market maker might be exposed to the rising options prices that occur as volatility increases. However, instruments such VIX futures and calls on VIX futures changed all that, as a short position in options could then have its Vega reduced to almost nothing with the appropriate number of VIX contracts.

Traders use VIX futures and options to hedge in times of uncertain volatility in the SPX or the SPY.¹ When volatility traders notice a spread between the VIX and VIX futures, they realize that a correction is probable. Therefore, if VIX is trading higher than the VIX futures price, then buying calls in both SPY and VIX will make money because either a) VIX goes down and the SPY goes up which places the SPY call in-the-money, or b) the VIX stays up and the VIX futures close the gap which places the VIX call in-the-money. A similar strategy with puts on SPY and VIX can be devised when the VIX is significantly lower than the VIX futures price. The rule of thumb is: provided that the spread between VIX and VIX futures is wide enough, the correction in VIX will create enough change in the market that one of these straddles can cover its initial cost and provide some profit to the investor. The operative word in the strategy mentioned in this paragraph is *spread*, which is precisely what we are looking for as we filter for the risk-premium.

¹SPX is the S&P500; SPY is the tracking stock for the S&P 500.

Chapter 3

Stochastic Volatility Filter for Heston Model

With discrete observations, we derive a Bayesian stochastic volatility filter for a Heston model. Let Y_t denote the log-price of the equity, and let $\sqrt{X_t}$ denote volatility. The dynamics of the processes are,

$$dX_t = \kappa(\bar{X} - X_t)dt + \gamma\sqrt{X_t}dB_t \quad (3.1)$$

$$dY_t = \left(\mu - \frac{1}{2}X_t\right)dt + \sqrt{X_t}\left(\rho dB_t + \sqrt{1-\rho^2}dW_t\right) \quad (3.2)$$

and the interpretation of the model parameters is as follows:

- \bar{X} = the long-time average of X_t
- κ = the rate of mean-reversion (on X_t)
- γ = volatility of volatility
- ρ = models volatility leverage effect when $\rho \in (-1, 0)$
- μ = the mean-rate of returns on the stock

Certain restrictions on the parameters need to be put in place, such as the Feller condition: $\gamma \leq \sqrt{2\kappa\bar{X}}$ to insure the X_t is well-defined (see chapter volatility time scales in [24]). There are times $(t_n)_n$ for which the process is actually observed, and if our model is correct¹ then $\mathcal{F}_t^m = \mathcal{F}_t^Y = \sigma\{Y_{t_n} : t_n \leq t\}$.

¹‘correct’ not only means that the processes follow these parametric SDEs, but it also means that $\{W_t\}$ and $\{B_t\}$ are idiosyncratic noises that are endemic to this system and not correlated with other data in the market.

3.1 The Filter

We showed in previous lectures that X_t is measurable when Y is observed continuously. It is also straight forward to show that $\mathbb{E}[g(X_t)|\mathcal{F}_t^Y] \rightarrow g(X_t)$ as the partition of $[0, t]$ shrinks to zero.

Lemma 3.1.1. *For any $N \in \mathbb{Z}^+$ let there be a partition of $[0, t]$ into N -many points, and let \mathcal{F}_t^N denote the filtration generated by the observations of Y at the partitioned points (i.e. $\mathcal{F}_t^N = \sigma\{\{Y_{t_n}\}_{n=0}^N\}$). If the filtrations are increasing with N , then in the context of the model given by (3.1) and (3.2) we have*

$$\mathbb{E}[g(X_t)|\mathcal{F}_t^N] \rightarrow g(X_t) \quad a.s.$$

as $N \nearrow \infty$ for any function $g(x)$.

Proof. Let $\mathcal{F}_t^Y = \sigma\{Y_s : s \leq t\}$. We assume that the filtrations are increasing with N , and they are certainly bounded by \mathcal{F}_t^Y ,

$$\mathcal{F}_t^N \subset \mathcal{F}_t^{N+1} \subset \dots \subseteq \mathcal{F}_t^Y.$$

Therefore there exists \mathcal{F}_t^* such that $\bigvee_{N=1}^{\infty} \mathcal{F}_t^N = \mathcal{F}_t^*$. Now because the path of Y_t is continuous on all sets of non-zero probability, the information contained in \mathcal{F}_t^* is enough to measure the event $\{Y_t = y\}$ for any $y \in \mathbb{R}$ even if t is not a partition point for any finite N . Therefore, $\mathcal{F}_t^* = \mathcal{F}_t^Y$ a.s., and by the Lévy 0-1 law we have

$$\lim_N \mathbb{E}[g(X_t)|\mathcal{F}_t^N] = \mathbb{E}[g(X_t)|\mathcal{F}_t^*] = \mathbb{E}[g(X_t)|\mathcal{F}_t^Y] = g(X_t)$$

which proves the lemma. □

From lemma 3.1.1 we know that our posterior estimates are consistent as we refine the partition. Now we need to determine the filter for a specific partition. From here forward consider a specific finite partition of the time domain, and we only consider the filter at times when data has arrived. For ease in notation we let $X_n = X_{t_n}$, $Y_n = Y_{t_n}$ and $\mathcal{F}_n^Y = \mathcal{F}_{t_n}^Y$.

Proposition 3.1.1. *Let (Y_t, X_t) be the price and volatility processes in the Heston model from (3.1) and (3.2), and assume the Feller condition $\gamma^2 \leq 2\bar{X}\kappa$. Then there is a kernel that gives X 's transition density, $e^{Q^*\Delta t}(x|v) = \frac{d}{dx}\mathbb{P}(X_{t+\Delta t} \leq x|X_t = v)$ for any $x, v \in \mathbb{R}^+$, and the filtering distribution for X_n at observation time $t_n = n\Delta t$ has a density. This density is given recursively as*

$$\begin{aligned} & \pi_n(x) \\ &= \frac{1}{c_n} \int \mathbb{E} \left[\mathbb{L}(y | (X_u)_{\{t_{n-1} \leq u \leq t_n\}}, Y_{n-1}) \middle| X_n = x, X_{n-1} = v, Y_{n-1} \right] e^{Q^*\Delta t}(x|v) \pi_{n-1}(dv) \bigg|_{y=Y_n}, \end{aligned} \tag{3.3}$$

for almost-everywhere $x \in \mathbb{R}^+$, where c_n is a normalizing constant, and \mathbb{L} is the likelihood of the any path $(x_u)_{\{t_{n-1} \leq u \leq t_n\}}$ given observations Y_n and Y_{n-1} , and is given by

$$\mathbb{L}(y|(x_u)_{\{t_{n-1} \leq u \leq t_n\}}, Y_{n-1}) = \frac{\exp \left\{ -\frac{1}{2} \left(\frac{(y - Y_{n-1}) - \left(\mu \Delta t - \int_{t_{n-1}}^{t_n} x_u du \right) - \rho \xi_n(x)}{\sqrt{(1 - \rho^2) \int_{t_{n-1}}^{t_n} x_u du}} \right)^2 \right\}}{\sqrt{(1 - \rho^2) \int_{t_{n-1}}^{t_n} x_u du}}$$

with

$$\xi_n(x) = \frac{1}{\gamma} \left\{ \Delta x_{n-1} - \kappa \left(\bar{X} \Delta t - \int_{t_{n-1}}^{t_n} x_u du \right) \right\}.$$

Proof. Given the Feller condition, the CIR process $dX_t = \kappa(\bar{X} - X_t)dt + \gamma\sqrt{X_t}dB_t$ is well-known to have a transition density that can be written in terms of a modified Bessel function (see [1]), and so $\Gamma^{\Delta t}(\cdot|v)$ is a smooth density function for all $v \geq 0$. Furthermore, it was shown in [20] that (Y_t, X_t) has a smooth transition density function, that is,

$$P^{\Delta t}(y, x|s, v) \doteq \frac{\partial^2}{\partial y \partial x} \mathbb{P}(Y_n \leq y, X_n \leq x | Y_{n-1} = s, X_{n-1} = v)$$

is smooth for $x > 0$, $y > 0$, and $\Delta t > 0$, and does not collect mass at $x = 0$ or $y = 0$. Hence, the filter has a density that can be written using Bayes rule:

$$\pi_n(x) = \frac{\int P^{\Delta t}(Y_n, x | Y_{n-1}, v) \pi_{n-1}(v) dv}{\int [\text{numerator}] dx},$$

where we don't need to assume smoothness of π_{n-1} because it is smoothed by its convolution with $P^{\Delta t}$ in the dv -integral.

Now, from equation (3.2) we notice the following:

$$\begin{aligned} Y_n - Y_{n-1} &= \mu \Delta t - \frac{1}{2} \int_{t_{n-1}}^{t_n} X_u du + \rho \int_{t_{n-1}}^{t_n} \sqrt{X_u} dB_u + \sqrt{1 - \rho^2} \int_{t_{n-1}}^{t_n} \sqrt{X_u} dW_u \\ &= {}_d \mu \Delta t - \frac{1}{2} \int_{t_{n-1}}^{t_n} X_u du + \rho \int_{t_{n-1}}^{t_n} \sqrt{X_u} dB_u + \sqrt{(1 - \rho^2) \int_{t_{n-1}}^{t_n} X_u du} \mathcal{Z} \end{aligned}$$

where “ $=_d$ ” signifies equivalence in distribution, and \mathcal{Z} is another independent standard normal random variable. This means that conditional on the path $(X_u)_{t_{n-1} \leq u \leq t_n}$ and S_{n-1} ,

$$\frac{Y_n - Y_{n-1} - \left(\mu \Delta t - \frac{1}{2} \int_{t_{n-1}}^{t_n} X_u du + \rho \int_{t_{n-1}}^{t_n} \sqrt{X_u} dB_u \right)}{\sqrt{(1 - \rho^2) \int_{t_{n-1}}^{t_n} X_u du}} = {}_d \mathcal{Z}.$$

Then noticing ξ_n evaluated at $(X_u)_{t_{n-1} \leq u \leq t_n}$ is the the same as $\xi_n(X) = \int_{t_{n-1}}^{t_n} \sqrt{X_u} dB_u$, it follows that

$$\frac{Y_n - Y_{n-1} - \left(\mu \Delta t - \frac{1}{2} \int_{t_{n-1}}^{t_n} X_u du + \rho \xi_n(X) \right)}{\sqrt{(1 - \rho^2) \int_{t_{n-1}}^{t_n} X_u du}} =_d \mathcal{Z}.$$

This shows the likelihood of the path $(X_u)_{t_{n-1} \leq u \leq t_n}$ given S_{n-1} and $S_n = y$ is in fact the function \mathbb{L} .

Finally, given Bayes rule for the density π_n , the expression in equation (3.3) displays the filter using a probabilistic representation of the transition density:

$$\begin{aligned} & P^{\Delta t}(y, x | Y_{n-1}, v) \\ &= \frac{\partial}{\partial y} \int_0^y P^{\Delta t}(z, x | Y_{n-1}, v) dz \\ &= \frac{\partial}{\partial y} \mathbb{P}(Y_n \leq y | X_n = x, X_{n-1} = v, Y_{n-1}) \Gamma^{\Delta t}(x | v) \\ &= \frac{\partial}{\partial y} \mathbb{E} \left\{ \mathbf{1}_{Y_n \leq y} \middle| X_n = x, X_{n-1} = v, Y_{n-1} \right\} \Gamma^{\Delta t}(x | v) \\ &= \frac{\partial}{\partial y} \mathbb{E} \left[\mathbb{E} \left\{ \mathbf{1}_{Y_n \leq y} \middle| (X_u)_{\{t_{n-1} \leq u \leq t_n\}}, Y_{n-1} \right\} \middle| X_n = x, X_{n-1} = v, Y_n, Y_{n-1} \right] \Gamma^{\Delta t}(x | v) \\ &= \mathbb{E} \left[\frac{\partial}{\partial y} \mathbb{E} \left\{ \mathbf{1}_{Y_n \leq y} \middle| (X_u)_{\{t_{n-1} \leq u \leq t_n\}}, Y_{n-1} \right\} \middle| X_n = x, X_{n-1} = v, Y_n, Y_{n-1} \right] \Gamma^{\Delta t}(x | v) \\ &\propto \mathbb{E} \left[\mathbb{L}(y | (X_u)_{\{t_{n-1} \leq u \leq t_n\}}, Y_{n-1}) \middle| X_n = x, X_{n-1} = v, Y_{n-1} \right] e^{Q^* \Delta t}(x | v). \end{aligned}$$

Lastly, when computing the likelihood based on the time- n observation, the last line is evaluated at $y = Y_n$. This completes the proof of the proposition. \square

At this point it seems that the filter is rather complicated, and would be difficult to implement in real-time. Often times, what one might do is consider a discrete scheme that approximates the SDEs:

$$\Delta Y_{n-1} = (\mu - .5 X_{n-1}) \Delta t + \sqrt{X_{n-1}} (\rho \Delta B_{n-1} + \sqrt{1 - \rho^2} \Delta W_{k-1}) \quad (3.4)$$

$$\Delta X_{n-1} = \kappa (\bar{X} - X_{n-1}) \Delta t + \gamma \sqrt{X_{n-1}} \Delta B_{n-1} \quad (3.5)$$

where $\Delta Y_{n-1} = Y_n - Y_{n-1}$ and $\Delta X_{n-1} = X_n - X_{n-1}$. Using (3.4) and (3.5), we can impute an approximate filtering density to the density given in Proposition 3.1.1:

$$\tilde{\pi}_n(x) = \frac{1}{c_n} \int \psi_n(x, v) e^{Q^* \Delta t}(x | v) \tilde{\pi}_{n-1}(dv) \quad (3.6)$$

where $\psi_n(x, v)$ is the likelihood of $\{X_n = x, X_{n-1} = v\}$ given $\{Y_n, Y_{n-1}\}$,

$$\psi_n(x, v) = \frac{1}{\sqrt{v(1 - \rho^2)\Delta t}} \exp \left\{ -\frac{(\Delta Y_{n-1} - (\mu - .5v)\Delta t - \sqrt{v}\rho\Delta B_{n-1}(x, v))^2}{2v(1 - \rho^2)\Delta t} \right\}$$

$$\text{with} \quad \Delta B_{n-1}(x, v) = \frac{1}{\gamma\sqrt{v}} (x - v - \kappa(\bar{X} - v)\Delta t).$$

If the model simplification given by (3.4) and (3.5) can be considered ‘correct’, then there is no need to dispute the validity of the filter given by (3.6). But in general, if one knows apriori that the continuous-time SDEs are the correct model, then there needs to be some analysis to verify that the approximate filter (such as that in (3.6)) converges as $\Delta t \searrow 0$, that is

$$\left| \int g(x) \tilde{\pi}_n(dx) - \int g(x) \pi_n(dx) \right| \rightarrow 0$$

in a strong sense as $\Delta t \rightarrow 0$. As was mentioned earlier, it is well-known (see Kushner [33]) that approximate filters that are sometimes consistent, but the results in [33] do not apply to the Heston model.

3.2 Extracting the Risk-Premium

Under the physical measure, we have,

$$\mathbb{E}X_t = \mathbb{E}X_0 e^{-\kappa t} + \bar{X}(1 - e^{-\kappa t})$$

and so the expected value of realized variance ($RV_{[0,T]} \doteq \frac{1}{T} \int_0^T X_s ds$) is

$$\mathbb{E}_0 RV_{[0,T]} = \bar{X} - \frac{\bar{X} - \mathbb{E}_0 X_0}{\kappa T} (1 - e^{-\kappa T})$$

where we have (without loss of generality) considered the case at time 0, and we have denoted the posterior expectation as $\mathbb{E}_0^*[\cdot] = \mathbb{E}^*[\cdot | \mathcal{F}_0^Y]$ and $\mathbb{E}_0[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_0^Y]$. Therefore, there is the following close-formula for the risk-premium

$$\begin{aligned} \mathcal{RP}_0 &\doteq \mathbb{E}_0^* RV_{[0,T]} - \mathbb{E}_0 RV_{[0,T]} \\ &= \mathbb{E}_0^* RV_{[0,T]} - \bar{X} + \frac{\bar{X} - \mathbb{E}_0 X_0}{\kappa T} (1 - e^{-\kappa T}). \end{aligned}$$

This expression for the risk-premium holds whenever volatility-squared is modeled with a mean-reverting SDE with drift term $\kappa(\bar{X} - X_t)$, not just the Heston model.

Under the risk-neutral measure, the market adds a risk-premium term to dX_t :

$$dX_t = \kappa(\bar{X} - X_t)dt - \Lambda_t X_t dt + \gamma \sqrt{X_t} dB_t^*$$

where B_t^* Brownian motion under the market measure, and Λ_t is the market price of volatility risk. We leave the modeling of Λ_t open here because we will not delve deeply into its correlation structure. However, Λ_t is most likely thought of as a mean-reverting process and could be modeled as such. Under the market's measure there is the following expectation of variance

$$\mathbb{E}_0^* X_t = \mathbb{E}_0 X_t - \int_0^t \mathbb{E}_0^*[X_s \Lambda_s] e^{-\kappa(t-s)} ds.$$

From this, we see that $\mathbb{E}_0^* RV_{[0,T]}$ can be written as follows,

$$\begin{aligned} \mathbb{E}_0^* RV_{[0,T]} &= \frac{1}{T} \int_0^T \mathbb{E}_0^* X_t dt = \frac{1}{T} \int_0^T \mathbb{E}_0 X_t ds - \frac{1}{T} \int_0^T \int_0^t \mathbb{E}_0^*[X_s \Lambda_s] e^{-\kappa(t-s)} ds dt \\ &= \mathbb{E}_0[RV_{0,T}] - \underbrace{\frac{1}{\kappa T} \int_0^T \left(1 - e^{-\kappa(T-s)}\right) \mathbb{E}_0^*[X_s \Lambda_s] ds}_{\text{risk-premium}} \quad (**). \end{aligned}$$

In (**), notice that if X and Λ are independent, then for $\kappa \gg 1$ the risk-premium simplifies to

$$\frac{1}{\kappa T} \int_0^T \left(1 - e^{-\kappa(T-s)}\right) \mathbb{E}_0^*[X_s \Lambda_s] ds \sim -\frac{\bar{X}}{T} \mathbb{E}_0^* \Lambda_T.$$

3.3 Filtering Average Volatility in Fast Time-Scales

The purpose of fast time-scales in volatility modeling is to capture mean reverting effects that occur on the order of 2 to 3 days. In the Heston model, suppose we are in a fast time-scale where $\gamma \sim \sqrt{\kappa}$ for κ large. Then, it can be shown that the distribution of X_t settles into a Γ distribution almost instantaneously,

$$X_t \Rightarrow \Gamma\left(2\bar{X}, \frac{1}{2}\right) \quad \text{as } \kappa \nearrow \infty$$

for all $t > 0$. Therefore, there is a fast-averaging of the realized variance,

$$\frac{1}{T} \int_0^T X_s ds \rightarrow \bar{X} \quad \text{in probability as } \kappa \nearrow \infty$$

and so the expected payoff of any contract that is a function of realized variance will be deterministic unless \bar{X} is random and/or unknown. Thus, building a risk-premium into the dynamics of X_t in the manner that we did in the previous section will not be meaningful in fast time-scales. An alternative idea would be to take \bar{X} as a hidden regime-process that is

governed by another Markov chain, and thus adds another dimension to the HMM. Then, we can take realized variance as our observations and write a filter to estimate the regime \bar{X}_t . We do this as follows:

Take \bar{X}_t to be a Markov chain with generator Q , for which we assume the standard structure for changes; changes in \bar{X}_t are governed by a Poisson jump process so that over a time interval of length Δt , the probability of \bar{X}_t changing states more than once is $o(\Delta t)$. The new dynamics of X_t are then

$$dX_t = \kappa(\bar{X}_t - X_t)dt + \gamma\sqrt{X_t}dB_t$$

and for such a model the realized variance is a random variable in fast time-scales,

$$\frac{1}{T} \int_0^T X_s ds \sim \frac{1}{T} \int_0^T \bar{X}_s ds \quad \text{for } \kappa \text{ large.}$$

Let $(t_{n,\ell})_{n,\ell}$ be a partition of some finite time interval, say 10 years, where n denotes the n th week, and ℓ denotes the ℓ th trade. Then

$$t_{n+1,\ell} - t_{n,\ell} = \Delta t = 1 \text{ week}, \quad \text{for any } \ell$$

$$t_{n,\ell+1} - t_{n,\ell} = \text{time from } \ell\text{th quote until the next trade during week } n.$$

We have observations on Y_t at each $t_{n,\ell}$ (i.e. $Y_{n,\ell} = Y_{t_{n,\ell}}$ is the ℓ th observations on the n th week). For any n and ℓ let $\Delta Y_{n,\ell} = Y_{n,\ell+1} - Y_{n,\ell}$, and for simplicity let $t_n = t_{n,0}$. We then have the following model for weekly observations on realized variance,

$$Z_{n+1} \doteq \frac{1}{\Delta t} \sum_{\ell} (\Delta Y_{n,\ell})^2 = \frac{1}{\Delta t} \int_{t_n}^{t_{n+1}} \bar{X}_s ds + \epsilon_{n+1}$$

where ϵ_{n+1} is a noise process with $\mathbb{E}\epsilon_n\epsilon_{n'} = 0$ if $n \neq n'$. Clearly, Z_n is observable and (Z_n, \bar{X}_n) is a Markov process. From here the goal is to estimate the the state-space and transition rates of \bar{X}_t , and then apply the nonlinear filtering results in estimating the variance risk-premium. Letting $\bar{\pi}_n(x_i) = \mathbb{P}(\bar{X}_{t_n} = x_i | Y_{0:n})$, we have an estimate of the physical measure's expectation of realized variance:

$$\mathbb{E}_{t_n} \left[\frac{1}{T} \int_{t_n}^{t_n+T} X_s ds \right] \approx \mathbb{E}_{t_n} \left[\frac{1}{T} \int_{t_n}^{t_n+T} \bar{X}_s ds \right] = \frac{1}{T} \int_{t_n}^{t_n+T} \sum_i x_i e^{Q^*(s-t_n)} \bar{\pi}_n(x_i) ds ,$$

for $\kappa \sim \gamma^2 \gg 1$.

Chapter 4

The Zakai Equation

Let $X_t \in \mathcal{S}$ be a Markov process with generator Q . Let the domain of Q be denoted by $\mathcal{B}(Q)$, and let $\mathcal{B}_b(Q)$ denote the subset of bounded functions in $\mathcal{B}(Q)$. For any function $g(x) \in \mathcal{B}_b(Q)$ we have the following limit:

$$\frac{\mathbb{E}[g(X_{t+\Delta t})|X_t = x] - g(x)}{\Delta t} \rightarrow Qg(x)$$

as $\Delta t \searrow 0$, for any $x \in \mathcal{S}$. If Q is densely-defined and its resolvent set includes all positive real numbers, then the Hille-Yosida theorem applies, allowing us to write the the distribution of X_t with a contraction semi-group. In these notes we assume that such conditions hold and that the transition density/mass is generated by an operator semigroup denoted by e^{Q^*t} .

A standard nonlinear filtering problem in SDE theory assumes that X_t is unobserved and that a process Y_t is given by an SDE

$$dY_t = h(t, X_t)dt + \gamma dW_t \quad \text{observed} \quad (4.1)$$

where W_t is an independent Wiener process, $\gamma > 0$ and we assume that $h(t, \cdot)$ is bounded for all $t < \infty$.

The pair (X_t, Y_t) is an HMM for which filtering can be used to find the posterior distribution. Let $\mathcal{F}_t^Y = \sigma\{Y_s : s \leq t\}$, and for any measurable function $g(x)$ let

$$\hat{g}_t = \mathbb{E}[g(X_t)|\mathcal{F}_t^Y].$$

The posterior expectation of \hat{g}_t is ultimately what is desired from filtering, but the methods for obtaining the posterior distribution are quite involved. The discrete Bayesian tools that we've used in earlier lectures cannot be used here because we are in a continuum that does not allow us to break apart the layers of the HMM. Instead, we will exploit well-known ideas from SDE theory to obtain the differentials for the filtering distribution. In particular, we will use the Girsanov theorem to obtain the Zakai equation.

4.1 Discrete Motivation from Bayesian Perspective

In their book, Karatzas and Shreve [30] give a discrete motivation for how the Girsanov theorem works. In a similar fashion, we consider a discrete problem and then construct a change of measure from the ratio of the appropriate densities. We then show how it is analogous to its continuous-time counterpart, and leads to a discrete approximation of the Zakai equation.

To do so, we start by considering a discrete-time analogue of (6.1),

$$Y_n = Y_{n-1} + h(t_{n-1}, X_{n-1})\Delta t + \gamma\Delta W_{n-1},$$

and assume that the unconditional distribution of X_t is a density for all $t \geq 0$ (the same idea will be applicable when X_t 's distribution has a mass function). We can easily apply Bayes theorem to obtain the filtering distribution on a Borel set \mathcal{A} :

$$\pi_n(\mathcal{A}) = \frac{1}{c_n} \int \psi_n(v) e^{Q^* \Delta t}(\mathcal{A}|v) \pi_{n-1}(dv)$$

where $e^{Q^* \Delta t}$ represents kernel of X_t 's transition densities, the likelihood function is

$$\psi_n(v) = \exp \left\{ -.5 \left(\frac{\Delta Y_{n-1} - h(t_{n-1}, v)\Delta t}{\gamma\sqrt{\Delta t}} \right)^2 \right\}, \quad \text{with } \Delta Y_{n-1} = Y_n - Y_{n-1},$$

and c_n is a normalizing constant. The Lebesgue differentiation theorem can be applied to obtain the density of the posterior,

$$\frac{1}{|\mathcal{A}|} \pi_n(\mathcal{A}) \rightarrow \pi_n(dx) = \frac{1}{c_n} \int \psi_n(v) e^{Q^* \Delta t}(dx|v) \pi_{n-1}(dv)$$

when \mathcal{A} shrinks nicely to $\{x\}$.

Keeping this discrete model and filter in mind, let's shift our attention to a joint density function of all observations and a possible path (x_0, x_1, \dots, x_n) taken by (X_0, X_1, \dots, X_n) ,

$$\begin{aligned} dp_n &\doteq \mathbb{P}(Y_0, Y_1, \dots, Y_n; dx_0, dx_1, \dots, dx_n) \\ &= \mathbb{P}(Y_0, Y_1, \dots, Y_n | x_0, x_1, \dots, x_n) \mathbb{P}(dx_0, dx_1, \dots, dx_n) \\ &= \left(\prod_{\ell=0}^{n-1} \psi_\ell(x_\ell) \right) \times \mathbb{P}(dx_0, dx_1, \dots, dx_n) = \left(\prod_{\ell=0}^{n-1} e^{-.5 \left(\frac{\Delta Y_\ell - h(t_\ell, x_\ell)\Delta t}{\gamma\sqrt{\Delta t}} \right)^2} \right) \times \mathbb{P}(dx_0, dx_1, \dots, dx_n) \\ &= e^{-.5 \sum_{\ell=0}^{n-1} \left(\frac{\Delta Y_\ell - h(t_\ell, x_\ell)\Delta t}{\gamma\sqrt{\Delta t}} \right)^2} \times \mathbb{P}(dx_0, dx_1, \dots, dx_n) \end{aligned}$$

where $\mathbb{P}(dx_0, dx_1, \dots, dx_n)$ can be obtained using the exponential of Q^* .

Next, consider an equivalent measure in $\tilde{\mathbb{P}}$ where $\Delta Y_n/\gamma$ is Brownian motion independent of X_n , and the law of X_n remains the same. Under this new measure the joint density function is

$$\begin{aligned} d\tilde{p}_n &\doteq \tilde{\mathbb{P}}(Y_0, Y_1, \dots, Y_n; dx_0, dx_1, \dots, dx_n) = \tilde{\mathbb{P}}(Y_0, Y_1, \dots, Y_n) \tilde{\mathbb{P}}(dx_0, dx_1, \dots, dx_n) \\ &= \left(\prod_{\ell=0}^{n-1} e^{-0.5 \left(\frac{\Delta Y_\ell}{\gamma \sqrt{\Delta t}} \right)^2} \right) \times \mathbb{P}(dx_0, dx_1, \dots, dx_n) \\ &= e^{-0.5 \sum_{\ell=0}^{n-1} \left(\frac{\Delta Y_\ell}{\gamma \sqrt{\Delta t}} \right)^2} \times \mathbb{P}(dx_0, dx_1, \dots, dx_n). \end{aligned}$$

The ratio of these densities is written follows:

$$\frac{dp_n}{d\tilde{p}_n} \doteq M_n \Big|_{\mathcal{F}_n^Y} = \exp \left\{ \sum_{\ell=0}^{n-1} \frac{h(t_\ell, x_\ell) \Delta Y_\ell}{\gamma^2} - \frac{1}{2} \sum_{\ell=0}^{n-1} \frac{h^2(t_\ell, x_\ell) \Delta t}{\gamma^2} \right\}$$

which is the likelihood ratio of any path for the discrete observation model. It is also the discrete analog of the exponential martingale that we use in the Girsanov theorem. Furthermore, we can use M_n to rewrite the filtering expectation in terms of the alternative measure,

$$\mathbb{E}[g(X_n) | \mathcal{F}_n^Y] = \frac{\int g(x_n) dp_n}{\int dp_n} = \frac{\int g(x_n) M_n d\tilde{p}_n}{\int M_n d\tilde{p}_n} = \frac{\tilde{\mathbb{E}}[g(X_n) M_n | \mathcal{F}_n^Y]}{\tilde{\mathbb{E}}[M_n | \mathcal{F}_n^Y]},$$

and if we define $\phi_n[g] = \tilde{\mathbb{E}}[g(X_n) M_n | \mathcal{F}_n^Y]$ we can write the filtering expectation as

$$\hat{g}_n = \mathbb{E}[g(X_n) | \mathcal{F}_n^Y] = \frac{\phi_n[g]}{\phi_n[1]}$$

for any function $g(x) \in \mathcal{B}_b(Q)$.

It turns out to be advantageous to analyze under $\tilde{\mathbb{P}}$ -measure because the dynamics ϕ_n are linear when we move to a continuum of observations. To get a sense of the linearity, consider the following discrete expansion for small Δt ,

$$\begin{aligned} \phi_{n+1}[g] &= \tilde{\mathbb{E}}[g(X_{n+1}) M_{n+1} | \mathcal{F}_{n+1}^Y] \\ &= \tilde{\mathbb{E}} \left[g(X_{n+1}) \left(1 + \frac{h(t_n, X_n)}{\gamma^2} \Delta Y_n \right) M_n \Big| \mathcal{F}_{n+1}^Y \right] + o(\Delta t) \\ &= \tilde{\mathbb{E}} \left[g(X_{n+1}) M_n \Big| \mathcal{F}_{n+1}^Y \right] + \tilde{\mathbb{E}} \left[g(X_{n+1}) \left(\frac{h(t_n, X_n)}{\gamma^2} \right) M_n \Big| \mathcal{F}_{n+1}^Y \right] \Delta Y_n + o(\Delta t) \end{aligned}$$

because $M_{n+1} = \left(1 + \frac{h(t_n, X_n)}{\gamma^2} \Delta Y_n\right) M_n + o(\Delta t)$ by a discrete interpretation of Itô's lemma, and ΔY_n is \mathcal{F}_{n+1}^Y -measurable. Now, because X is independent of Y under $\tilde{\mathbb{P}}$, the conditioning up to time $n+1$ is superfluous and we can reduce it down to time n ,

$$= \tilde{\mathbb{E}} \left[g(X_{n+1}) M_n \middle| \mathcal{F}_n^Y \right] + \tilde{\mathbb{E}} \left[g(X_{n+1}) \left(\frac{h(t_n, X_n)}{\gamma^2} \right) M_n \middle| \mathcal{F}_n^Y \right] \Delta Y_n + o(\Delta t).$$

Then we can again exploit the independence of X from Y under $\tilde{\mathbb{P}}$ and use an approximation of X 's backwards operator.

$$\begin{aligned} &= \tilde{\mathbb{E}} \left\{ \tilde{\mathbb{E}} \left[g(X_{n+1}) M_n \middle| \mathcal{F}_n^Y \vee X_n \right] \middle| \mathcal{F}_n^Y \right\} \\ &\quad + \tilde{\mathbb{E}} \left\{ \tilde{\mathbb{E}} \left[g(X_{n+1}) \left(\frac{h(t_n, X_n)}{\gamma^2} \right) M_n \middle| \mathcal{F}_n^Y \vee X_n \right] \middle| \mathcal{F}_n^Y \right\} \Delta Y_n + o(\Delta t) \\ &= \tilde{\mathbb{E}} \left\{ (I + Q\Delta t) \tilde{\mathbb{E}} \left[g(X_n) M_n \middle| \mathcal{F}_n^Y \vee X_n \right] \middle| \mathcal{F}_n^Y \right\} \\ &\quad + \tilde{\mathbb{E}} \left\{ (I + Q\Delta t) \tilde{\mathbb{E}} \left[g(X_n) \left(\frac{h(t_n, X_n)}{\gamma^2} \right) M_n \middle| \mathcal{F}_n^Y \vee X_n \right] \middle| \mathcal{F}_n^Y \right\} \Delta Y_n + o(\Delta t) \\ &= \tilde{\mathbb{E}} \left[(I + Q\Delta t) g(X_n) M_n \middle| \mathcal{F}_n^Y \right] + \tilde{\mathbb{E}} \left[(I + Q\Delta t) g(X_n) \left(\frac{h(t_n, X_n)}{\gamma^2} \right) M_n \middle| \mathcal{F}_n^Y \right] \Delta Y_n + o(\Delta t) \end{aligned}$$

and then using the fact that $\Delta Y_n \cdot \Delta t = o(\Delta t)$, we have

$$\begin{aligned} &= \tilde{\mathbb{E}}[g(X_n) M_n | \mathcal{F}_n^Y] + \tilde{\mathbb{E}}[Qg(X_n) M_n | \mathcal{F}_n^Y] \Delta t + \tilde{\mathbb{E}} \left[g(X_n) \frac{h(t_n, X_n)}{\gamma^2} M_n \middle| \mathcal{F}_n^Y \right] \Delta Y_n + o(\Delta t) \\ &= \phi_n[g] + \phi[Qg] \Delta t + \phi \left[g \frac{h}{\gamma^2} \right] \Delta Y_n + o(\Delta t) \end{aligned}$$

which foreshadows the Zakai equation in a discrete setting,

$$\Delta \phi_n[g] = \phi[Qg] \Delta t + \phi \left[g \frac{h}{\gamma^2} \right] \Delta Y_n + o(\Delta t).$$

4.2 Derivation of the Zakai Equation

In this section we use the Girsanov theorem as the main tool in a formal derivation of the nonlinear filtering equations in continuous time. For ease in notation we let $h(t, x) = h(x)$, but this does change the results because they would merely need to be rewritten to include the time dependence in $h(\cdot)$.

We start by considering the finite interval $[0, T]$ and defining the following exponential,

$$\begin{aligned} M_t &\doteq \exp \left\{ \frac{1}{2\gamma^2} \int_0^t h^2(X_s) ds + \frac{1}{\gamma} \int_0^t h(X_s) dW_s \right\} \\ &= \exp \left\{ -\frac{1}{2\gamma^2} \int_0^t h^2(X_s) ds + \frac{1}{\gamma^2} \int_0^t h(X_s) dY_s \right\} \end{aligned}$$

for all $t \in [0, T]$. Since it was initially assumed that W_t was independent, we can define an equivalent measure $\tilde{\mathbb{P}}$ by

$$d\tilde{\mathbb{P}} = M_T^{-1} d\mathbb{P}.$$

By the Girsanov theorem we know that

1. $\tilde{\mathbb{P}}$ is a probability measure, and
2. Y_t/γ is $\tilde{\mathbb{P}}$ -Brownian motion for $t \in [0, T]$, conditioned on X .

It is easy to show with moment generating functions that $\tilde{\mathbb{P}}(X_t \leq x) = \mathbb{P}(X_t \leq x)$. We can also easily show that M_t is a true $\tilde{\mathbb{P}}$ -martingale,

$$\begin{aligned} \tilde{\mathbb{E}} M_t &= \int M_t(\omega) d\tilde{\mathbb{P}}(\omega) = \int M_t(\omega) M_T^{-1}(\omega) d\mathbb{P}(\omega) = \mathbb{E}[M_t/M_T] \\ &= \mathbb{E} \left[\exp \left\{ -\frac{1}{2\gamma^2} \int_t^T h^2(X_s) ds - \frac{1}{\gamma} \int_t^T h(X_s) dW_s \right\} \right] \\ &= 1. \end{aligned}$$

Lastly, a simple lemma shows that X and Y are path-wise independent under $\tilde{\mathbb{P}}$:

Lemma 4.2.1. *X and Y are path-wise independent under $\tilde{\mathbb{P}}$.*

Proof. For an arbitrary path-wise function f_1 we have

$$\tilde{\mathbb{E}} f_1(Y/\gamma) = \tilde{\mathbb{E}}[\tilde{\mathbb{E}}[f_1(Y/\gamma)|X]] = \tilde{\mathbb{E}}[\mathbb{E}f_1(W)] = \mathbb{E}f_1(W)$$

showing that Y/γ is $\tilde{\mathbb{P}}$ -Brownian motion unconditional on X . Then for another arbitrary path-wise function f_2 we have

$$\begin{aligned} \tilde{\mathbb{E}} f_1(Y/\gamma) f_2(X) &= \tilde{\mathbb{E}}[f_2(X) \tilde{\mathbb{E}}[f_1(Y/\gamma)|X]] = \tilde{\mathbb{E}}[f_2(X) \mathbb{E}f_1(W)] \\ &= \mathbb{E}f_1(W) \tilde{\mathbb{E}}f_2(X) = \tilde{\mathbb{E}}f_1(Y/\gamma) \tilde{\mathbb{E}}f_2(X) \end{aligned}$$

and so X and Y are $\tilde{\mathbb{P}}$ -independent. □

From here forward, define the measure ϕ_t on $\mathcal{B}_b(Q)$ as

$$\phi_t[g] \doteq \tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_t^Y].$$

With this new measure we can express another important result regarding the Girsanov change of measure, namely the Kallianpur-Streibel formula:

Lemma 4.2.2. *Kallianpur-Streibel Formula:*

$$\mathbb{E}[g(X_t)|\mathcal{F}_t^Y] = \frac{\phi_t[g]}{\phi_t[1]}$$

for any $g(x) \in \mathcal{B}(Q)$.

Proof. For any $A \in \mathcal{F}_t^Y$ we have

$$\begin{aligned} \mathbb{E}\{\mathbf{1}_A \mathbb{E}[g(X_t)|\mathcal{F}_t^Y]\} &= \mathbb{E}[\mathbf{1}_A g(X_t)] = \tilde{\mathbb{E}}[\mathbf{1}_A g(X_t)M_t] \\ &= \tilde{\mathbb{E}}\left\{\mathbf{1}_A \tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_t^Y]\right\} = \tilde{\mathbb{E}}\left\{\mathbf{1}_A \frac{\tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_t^Y]}{\tilde{\mathbb{E}}[M_t|\mathcal{F}_t^Y]} \tilde{\mathbb{E}}[M_t|\mathcal{F}_t^Y]\right\} = \tilde{\mathbb{E}}\left\{\mathbf{1}_A \frac{\phi_t[g]}{\phi_t[1]} \tilde{\mathbb{E}}[M_t|\mathcal{F}_t^Y]\right\} \\ &= \tilde{\mathbb{E}}\left\{\tilde{\mathbb{E}}\left[\mathbf{1}_A \frac{\phi_t[g]}{\phi_t[1]} M_t \middle| \mathcal{F}_t^Y\right]\right\} = \tilde{\mathbb{E}}\left[\mathbf{1}_A \frac{\phi_t[g]}{\phi_t[1]} M_t\right] = \mathbb{E}\left[\mathbf{1}_A \frac{\phi_t[g]}{\phi_t[1]}\right] \end{aligned}$$

and since A was an arbitrary set, this shows that the result holds wp1. \square

Now, for any $t \in [0, T]$, we use Fubini's theorem to bring the differential inside the expectation, and from there we apply Itô's lemma, which gives us the following differential,

$$d\tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_T^Y] = \tilde{\mathbb{E}}[Qg(X_t)M_t|\mathcal{F}_T^Y]dt + \tilde{\mathbb{E}}\left[g(X_t)\frac{h(X_t)}{\gamma^2}M_t \middle| \mathcal{F}_T^Y\right]dY_t$$

for any $g(x) \in \mathcal{B}_b(Q)$. From this we can construct the integrated form of the differential,

$$\begin{aligned} &\tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_T^Y] \\ &= \tilde{\mathbb{E}}[g(X_0)|\mathcal{F}_T^Y] + \int_0^t \tilde{\mathbb{E}}[Qg(X_s)M_s|\mathcal{F}_T^Y]ds + \int_0^t \tilde{\mathbb{E}}\left[g(X_s)\frac{h(X_s)}{\gamma^2}M_s \middle| \mathcal{F}_T^Y\right]dY_s \end{aligned}$$

and since X is independent of Y under the $\tilde{\mathbb{P}}$ -measure, we can reduce the filtrations from \mathcal{F}_T^Y to \mathcal{F}_s^Y for all $s \leq T$, giving us,

$$\tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_t^Y] = \mathbb{E}[g(X_0)] + \int_0^t \tilde{\mathbb{E}}[Qg(X_s)M_s|\mathcal{F}_s^Y]ds + \int_0^t \tilde{\mathbb{E}}\left[g(X_s)\frac{h(X_s)}{\gamma^2}M_s \middle| \mathcal{F}_s^Y\right]dY_s \quad (4.2)$$

for all $t \leq T$. Inserting $\phi_s[\cdot]$ in (4.2) wherever possible and then taking the differential with respect to t , we have **the Zakai equation**:

$$d\phi_t[g] = \phi_t[Qg]dt + \phi_t \left[g \frac{h}{\gamma^2} \right] dY_t \quad (4.3)$$

for all $t \leq T$. The Zakai equation can also be considered for general unbounded functions $g(x) \in \mathcal{B}(Q)$, but we have restricted ourselves to the bounded case in order to insure that $\phi_t[g]$ is finite almost surely. Existence of solutions to (4.3) is straight-forward because we have derived it by differentiating $\tilde{\mathbb{E}}[g(X_t)M_t|\mathcal{F}_t^Y]$. Uniqueness of measure-valued solutions to (4.3) has been shown by Kurtz and Ocone [32] using a filtered martingale problem, and by Rozovsky [42] using a Radon measure representation of ϕ_t .

4.2.1 The Adjoint Zakai Equation

Depending on the nature of the filtering problem, the unnormalized probability measure $\phi_t[\cdot]$ may have a density/mass function. For example, suppose that $X_t \in \mathbb{R}$ is a diffusion process satisfying the SDE

$$dX_t = a(X_t)dt + \sigma dB_t$$

where $B_t \perp W_t$. Then the generator is $Q = \mathcal{L} = \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} \cdot + a(x) \frac{\partial}{\partial x} \cdot$ and the Zakai equation is

$$\begin{aligned} d\phi_t[g] &= \phi_t[\mathcal{L}g]dt + \frac{1}{\gamma^2} \phi_t[gh]dY_t \\ &= \frac{\sigma^2}{2} \phi_t \left[\frac{\partial^2}{\partial x^2} g \right] dt + \phi_t \left[a \frac{\partial}{\partial x} g \right] dt + \frac{1}{\gamma^2} \phi_t[gh]dY_t \end{aligned}$$

for any bounded function $g(x)$ with a 2nd derivative. Depending on $a(x), \sigma$ and the initial conditions, $\tilde{\pi}_t$ may be a density so that

$$\phi_t[g] = \int g(x) \tilde{\pi}_t(x) dx,$$

and provided that certain regularity conditions are met, the adjoint of the Zakai equation gives us an SPDE for $\tilde{\pi}_t$,

$$d\tilde{\pi}_t(x) = \mathcal{L}^* \tilde{\pi}_t(x)dt + \frac{h(x)}{\gamma^2} \tilde{\pi}_t(x)dY_t \quad (4.4)$$

with the initial condition $\tilde{\pi}_0(x) = \mathbb{P}(X_0 \in dx)$, and with the adjoint operator given by $\mathcal{L}^* = \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} \cdot - \frac{\partial}{\partial x} (a(x) \cdot)$. Existence and uniqueness of such densities is beyond the scope of these notes. Readers who are interested in regularity of solutions should read the book by Pardoux [37].

In the case of filtering distributions that are composed of mass functions, the adjoint equation is similar. For instance, if $X_t \in \{x_1, \dots, x_m\}$ is a finite-state Markov chain with generator Q , the adjoint equation holds without any regularity conditions,

$$\begin{aligned} d\tilde{\pi}_t(x_i) &= Q^* \tilde{\pi}_t(x_i) dt + \frac{h(x_i)}{\gamma^2} \tilde{\pi}_t(x_i) dY_t \\ &= \sum_{j=1}^m Q_{ji} \tilde{\pi}_t(x_j) dt + \frac{h(x_i)}{\gamma^2} \tilde{\pi}_t(x_i) dY_t. \end{aligned}$$

General existence and uniqueness for $\tilde{\pi}_t$ in this discrete case was shown by Rozovsky [41].

4.2.2 Kushner-Stratonovich Equation

Using the Zakai equation of (4.3) and observing that our assumption that h is bounded implies $\mathbb{P}(\phi_t[1] < \infty) = 1$ for all $t \in [0, T]$, we can apply Itô's lemma to obtain

$$d\hat{g}_t = d\left(\frac{\phi_t[g]}{\phi_t[1]}\right) = \widehat{Q}g_t dt + \frac{\widehat{g}h_t - \hat{g}_t \hat{h}_t}{\gamma^2} (dY_t - \hat{h}_t dt). \quad (4.5)$$

However, it should be mentioned that (4.5) was originally obtain a few years before the Zakai equation using other methods. Under the appropriate regularity conditions, the Kushner-Stratonovich equation is the adjoint (4.5) and is the nonlinear equation for the filtering distribution,

$$d\pi_t(x) = Q^* \pi_t(x) dt + \frac{h(x) - \hat{h}_t}{\gamma^2} \pi_t(x) (dY_t - \hat{h}_t dt)$$

where $\pi_t(x)$ is either a density or a mass function (depending on the type of problem).

4.2.3 Smoothing

The smoothing filter has been derived in [10], but in the case of regularized processes where the adjoint Zakai equation holds. In this section we derive a similar results but for the general case of functions in $\mathcal{B}_b(\mathcal{S})$, and we'll also derive the backward SPDE for smoothing in the regularized case.

Consider the times τ and t such that $0 \leq \tau \leq t$. The filtering expectation of $g(X_\tau)$ is

$$\mathbb{E}[g(X_\tau)|\mathcal{F}_t^Y] = \frac{\tilde{\mathbb{E}}[M_t g(X_\tau)|\mathcal{F}_t^Y]}{\tilde{\mathbb{E}}[M_t|\mathcal{F}_t^Y]} = \frac{\tilde{\mathbb{E}}[M_t g(X_\tau)|\mathcal{F}_t^Y]}{\phi_t[1]}.$$

For τ fixed and for t increasing, the Zakai equation is

$$\begin{aligned} d\tilde{\mathbb{E}}[M_t g(X_\tau)|\mathcal{F}_t^Y] &= \frac{1}{\gamma^2} \tilde{\mathbb{E}}[M_t h(X_t) g(X_\tau)|\mathcal{F}_t^Y] dY_t \quad \text{for } t > \tau, \\ \tilde{\mathbb{E}}[M_\tau g(X_\tau)|\mathcal{F}_\tau^Y] &= \phi_\tau[g]. \end{aligned}$$

This equation can be solved to obtain the smoothing distribution, but lacks a differential formula for changes in τ . However, in the regular case there is a backward SPDE that will provide a differential for changes in τ . This backward SDE will provide improved efficiency for coding and analysis.

Suppose there is sufficient regularity so that $\tilde{\pi}_t(x) \doteq \phi_t[\delta_x]$ satisfies the adjoint Zakai equation (see equation (4.4)). Given \mathcal{F}_t^Y , the smoothing filter for any time $\tau \in [0, t]$ is

$$\begin{aligned}
\frac{d}{dx} \mathbb{P}(X_\tau \leq x | \mathcal{F}_t^Y) &= \frac{\tilde{\mathbb{E}}[M_t \delta_x(X_\tau) | \mathcal{F}_t^Y]}{\tilde{\mathbb{E}}[M_t | \mathcal{F}_t^Y]} \\
&= \frac{\tilde{\mathbb{E}}[M_\tau \delta_x(X_\tau) \tilde{\mathbb{E}}[M_t/M_\tau | \mathcal{F}_t^Y \vee \mathcal{F}_\tau^X] | \mathcal{F}_t^Y]}{\tilde{\mathbb{E}}[M_t | \mathcal{F}_t^Y]} \\
&= \frac{\tilde{\mathbb{E}}[M_\tau \delta_x(X_\tau) \tilde{\mathbb{E}}[M_t/M_\tau | \mathcal{F}_t^Y \vee \{X_\tau = x\}] | \mathcal{F}_t^Y]}{\tilde{\mathbb{E}}[M_t | \mathcal{F}_t^Y]} \\
&= \frac{\tilde{\mathbb{E}}[M_\tau \delta_x(X_\tau) | \mathcal{F}_t^Y] \cdot \tilde{\mathbb{E}}[M_t/M_\tau | \mathcal{F}_t^Y \vee \{X_\tau = x\}]}{\tilde{\mathbb{E}}[M_t | \mathcal{F}_t^Y]} \\
&= \frac{\tilde{\mathbb{E}}[M_\tau \delta_x(X_\tau) | \mathcal{F}_\tau^Y] \cdot \tilde{\mathbb{E}}[M_t/M_\tau | \mathcal{F}_t^Y \vee \{X_\tau = x\}]}{\tilde{\mathbb{E}}[M_t | \mathcal{F}_t^Y]} \\
&= \frac{\tilde{\pi}_\tau(x) \alpha_{\tau,t}(x)}{\phi_t[1]},
\end{aligned}$$

where $\alpha_{\tau,t}$ is define as

$$\begin{aligned}
\alpha_{\tau,t}(x) &\doteq \tilde{\mathbb{E}} \left[M_t/M_\tau \middle| \mathcal{F}_t^Y \vee \{X_\tau = x\} \right] \\
&= \tilde{\mathbb{E}} \left[\exp \left\{ -\frac{1}{2\gamma^2} \int_\tau^t h^2(X_s) ds + \frac{1}{\gamma^2} \int_\tau^t h(X_s) dY_s \right\} \middle| \mathcal{F}_t^Y \vee \{X_\tau = x\} \right].
\end{aligned}$$

The function $\alpha_{\tau,t}(x)$ is the smoothing component and satisfies the following backward SPDE

$$\begin{aligned}
d\alpha_{\tau,t}(x) &= Q\alpha_{\tau,t}(x)d\tau - \frac{1}{\gamma^2} h(x)\alpha_{\tau,t}(x)dY_\tau \quad \text{for } \tau \leq t, \\
\alpha_{t,t} &\equiv 1,
\end{aligned}$$

or in integrated form

$$\alpha_{\tau,t}(x) = 1 - \int_\tau^t Q\alpha_{s,t}(x)ds + \frac{1}{\gamma^2} \int_\tau^t h(x)\alpha_{s,t}(x)dY_s.$$

Chapter 5

The Innovations Approach

Let $X_t \in \mathcal{S}$ be a Markov process with generator Q . Let the domain of Q be denoted by $\mathcal{B}(Q)$, and let $\mathcal{B}_b(Q)$ denote the subset of bounded functions in $\mathcal{B}(Q)$. For any function $g(x) \in \mathcal{B}_b(Q)$ we have the following limit:

$$\frac{\mathbb{E}[g(X_{t+\Delta t})|X_t = x] - g(x)}{\Delta t} \rightarrow Qg(x)$$

as $\Delta t \searrow 0$, for any $x \in \mathcal{S}$. If Q is densely-defined and its resolvent set includes all positive real numbers, then the Hille-Yosida theorem applies, allowing us to write the the distribution of X_t with a contraction semi-group. In these notes we assume that such conditions hold and that the transition density/mass is generated by an operator semigroup denoted by e^{Q^*t} .

A standard nonlinear filtering problem in SDE theory assumes that X_t is unobserved and that a process Y_t is given by an SDE

$$dY_t = h(X_t)dt + \gamma dW_t \quad \text{observed} \quad (5.1)$$

where W_t is an independent Wiener process, $\gamma > 0$ and we assume that the h is bounded. Let the filtration $\mathcal{F}_t^Y = \sigma\{Y_s : s \leq t\}$ so that for an integrable function $g(x)$ we have

$$\hat{g}_t \doteq \mathbb{E}[g(X_t)|\mathcal{F}_t^Y]$$

for all $t \in [0, T]$.

5.1 Innovations Brownian Motion

Let ν_t denote the *innovations process* whose differential is given as follows

$$d\nu_t = dY_t - \hat{h}_t dt.$$

with $\nu_0 = 0$, and where $\hat{h}_t = \mathbb{E}[h(X_t)|\mathcal{F}_t^Y]$.

Proposition 5.1.1. *The process ν_t/γ is an \mathcal{F}_t^Y Brownian motion.*

Proof. It is clear that ν_t is (i) \mathcal{F}_t^Y -measurable, continuous and square integrable on $[0, T]$. To show that it is a local martingale, we take expectations for any $s \leq t$ as follows

$$\begin{aligned}
\mathbb{E}[\nu_t | \mathcal{F}_s^Y] - \nu_s &= \mathbb{E} \left[Y_t - \int_0^t \hat{h}_\tau d\tau \middle| \mathcal{F}_s^Y \right] - \left(Y_s - \int_0^s \hat{h}_\tau d\nu_\tau \right) \\
&= \mathbb{E} \left[\gamma(W_t - W_s) + \int_s^t h(X_\tau) d\tau \middle| \mathcal{F}_s^Y \right] - \mathbb{E} \left[\int_0^t \hat{h}_\tau d\tau \middle| \mathcal{F}_s^Y \right] + \int_0^s \hat{h}_\tau d\nu_\tau \\
&= \mathbb{E} \left[\gamma(W_t - W_s) + \int_s^t h(X_\tau) d\tau \middle| \mathcal{F}_s^Y \right] - \mathbb{E} \left[\int_s^t \hat{h}_\tau d\nu_\tau \middle| \mathcal{F}_s^Y \right] \\
&= \gamma \mathbb{E} \left[W_t - W_s \middle| \mathcal{F}_s^Y \right] + \mathbb{E} \left[\int_s^t h(X_\tau) d\tau - \int_s^t \hat{h}_\tau d\nu_\tau \middle| \mathcal{F}_s^Y \right] = 0.
\end{aligned}$$

Furthermore, the cross-variation of ν_t/γ is the same as the cross-variation of W_t , and so by the Lévy characterisation of Brownian motion (see Karatzas and Shreve [30]), ν_t/γ is also \mathcal{F}_t^Y Brownian motion. \square

Given that ν_t is \mathcal{F}_t^Y -Brownian motion, it may seem obvious that any L^2 -integrable and \mathcal{F}_T^Y measurable random variable has an integrated representation in terms of ν , but it is not easily seen that $\mathcal{F}_t^Y = \sigma\{\nu_s : s \leq t\}$. The following proposition provides a proof to verify that it is indeed true.

Proposition 5.1.2. *Every square integrable random variable N that is \mathcal{F}_T^Y -measurable, has a representation of the form*

$$N = \mathbb{E}N + \frac{1}{\gamma} \int_0^T f_s d\nu_s$$

where $f = \{f_s : s \leq T\}$ is progressively measurable and \mathcal{F}_t^Y -adapted and $\mathbb{E} \left[\int_0^T f_s^2 ds \right] < \infty$.

Proof. For all $t \in [0, T]$, define $Z_t \doteq \exp \left\{ -\frac{1}{\gamma^2} \int_0^t \hat{h}_s d\nu_s - \frac{1}{2\gamma^2} \int_0^t \hat{h}_s^2 ds \right\}$. Clearly, Z_t is an \mathcal{F}_t^Y -martingale, and so we can define an equivalent measure $\tilde{\mathbb{P}}$ with the following Radon-Nikodym derivative,

$$\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} = Z_T.$$

As a consequence of the Girsanov theorem, Y_t/γ is a $\tilde{\mathbb{P}}$ -Brownian motion. Then apply the martingale representation theorem,

$$Z_T^{-1}N = \tilde{\mathbb{E}}[Z_T^{-1}N] + \frac{1}{\gamma} \int_0^T q_s dY_s$$

$$= \tilde{\mathbb{E}}[Z_T^{-1}N] + \frac{1}{\gamma} \int_0^T q_s d\nu_s + \frac{1}{\gamma} \int_0^T q_s \hat{h}_s ds$$

where $q = \{q_s : s \leq T\}$ is adapted and $\tilde{\mathbb{P}}\left(\int_0^t q_s^2 < \infty\right) = \mathbb{P}\left(\int_0^t q_s^2 < \infty\right) = 1$. From here we can construct a $\tilde{\mathbb{P}}$ -martingale from $Z_T^{-1}N$,

$$\tilde{N}_t \doteq \tilde{\mathbb{E}}[Z_T^{-1}N | \mathcal{F}_t^Y]$$

and applying Itô's lemma to $\tilde{N}_t Z_t$ we have

$$\begin{aligned} d\left(\tilde{N}_t Z_t\right) &= -\frac{1}{\gamma^2} Z_t \tilde{N}_t \hat{h}_t d\nu_t + \frac{1}{\gamma} Z_t q_t d\nu_t + \frac{1}{\gamma} Z_t \hat{h}_t q_t dt - \frac{1}{\gamma} Z_t \hat{h}_t q_t dt \\ &= -\frac{1}{\gamma^2} Z_t \tilde{N}_t \hat{h}_t d\nu_t + \frac{1}{\gamma} Z_t q_t d\nu_t. \end{aligned}$$

Integrating from 0 to t , we have

$$\tilde{N}_t Z_t = \mathbb{E}N + \frac{1}{\gamma} \int_0^t \left(q_s - \frac{1}{\gamma} \tilde{N}_s \hat{h}_s\right) Z_s d\nu_s,$$

and therefore, setting $f_t = \left(q_t - \frac{1}{\gamma} \tilde{N}_t \hat{h}_t\right) Z_t$ for all $t \in [0, T]$, we have a unique \mathcal{F}_t^Y -adapted representation in terms of ν , and the proposition is proved. \square

Remark 3. For a more general proof of proposition 5.1.2, see proposition 2.31 on page 34 of Bain and Crisan [5]. They present the proof for generalized systems on unbounded time intervals, with merely the conditions that $\mathbb{E}\left[\int_0^t h^2(X_s) ds\right] < \infty$ and $\mathbb{P}\left(\int_0^t \hat{h}_s^2 ds < \infty\right) = 1$ for all $t < \infty$.

5.2 The Nonlinear Filter

In deriving the nonlinear filter with the innovations Brownian motion, it will be important to use the following martingale for any given $g \in \mathcal{B}_b(Q)$:

$$N_t \doteq \hat{g}_t - \int_0^t \widehat{Q} g_s ds$$

for all $t \in [0, T]$.

Lemma 5.2.1. N_t is an \mathcal{F}_t^Y -adapted martingale.

Proof. It suffices to show that $\mathbb{E}N_t = N_0$ for any $t \in [0, T]$, which we do as follows:

$$\begin{aligned}\mathbb{E}N_t &= \mathbb{E} \left[\hat{g}_t - \int_0^t \widehat{Q}g_s ds \right] = \mathbb{E}g_t - \int_0^t \mathbb{E} [\mathbb{E}[Qg(X_s)|\mathcal{F}_s]] ds \\ &= \mathbb{E}g_t - \int_0^t \mathbb{E}Qg(X_s)ds = \mathbb{E} \left[g_t - \int_0^t Qg(X_s)ds \right] = \mathbb{E}g(X_0) = N_0.\end{aligned}$$

□

Knowing that $\hat{g}_t - \int_0^t \widehat{Q}g_s ds$ is an \mathcal{F}_t^Y -martingale, we will apply proposition 5.1.2 as follows

$$\hat{g}_t - \int_0^t \widehat{Q}g_s ds = \hat{g}_0 + \frac{1}{\gamma} \int_0^t f_s d\nu_s$$

where f_t is an \mathcal{F}_t^Y -predictable process for any $t \in [0, T]$. From here, the main point in the derivation of the nonlinear filter is in finding the function f_t in terms of quantities that are more readily computable.

Theorem 5.2.1. The Nonlinear Filter. *For any function $g \in \mathcal{B}_b(Q)$, the nonlinear filter is given by the following SDE*

$$d\hat{g}_t = \widehat{Q}g_t dt + \frac{\widehat{g}h_t - \hat{g}_t \hat{h}_t}{\gamma^2} d\nu_t$$

for all $t \in [0, T]$.

Proof. We can apply proposition 5.1.2 to N_t and we get

$$N_t = \mathbb{E}N_0 + \frac{1}{\gamma} \int_0^t f_s d\nu_s = \hat{g}_0 + \frac{1}{\gamma} \int_0^t f_s d\nu_s,$$

thus defining the conditional expectation at time t as

$$\hat{g}_t = \hat{g}_0 + \int_0^t \widehat{Q}g_s ds + \frac{1}{\gamma} \int_0^t f_s d\nu_s.$$

From here, to complete the proof only requires us to identify f_t explicitly. For some process $\psi \in L^\infty[0, T]$, define ξ_t such that

$$d\xi_t = \frac{i}{\gamma} \xi_t \psi_t dY_t$$

with $\xi_0 = 1$. We then apply Itô's lemma to the following

$$d(\hat{g}_t \xi_t) = \widehat{Q}g_t \xi_t dt + \frac{1}{\gamma} f_t \xi_t d\nu_t + \frac{i}{\gamma} \hat{g}_t \xi_t \psi_t (d\nu_t + \hat{h}_t dt) + i \xi_t \psi_t f_t dt \quad (5.2)$$

$$d\mathbb{E}[g(X_t)\xi_t] = \mathbb{E}[Qg(X_t)\xi_t] dt + \frac{i}{\gamma} \mathbb{E}[g(X_t)\xi_t \psi_t h(X_t)] dt. \quad (5.3)$$

If we integrate the integrands in (5.2) and (5.3) from time 0 to time t , take expectations, multiply both sides by γ , and then subtract one from the other, and we are left with

$$\int_0^t i\psi_s \mathbb{E} \left[\xi_s \left(\gamma f_s - g(X_s)h(X_s) + \hat{g}_s \hat{h}_s \right) \right] ds = 0.$$

Hence, for almost every $t \in [0, T]$, we have

$$\mathbb{E} \left[\xi_t \left(\gamma f_t - g(X_t)h(X_t) + \hat{g}_t \hat{h}_t \right) \right] = 0 = \mathbb{E} \left[\xi_t \left(\gamma f_t - \mathbb{E}[g(X_t)h(X_t)|\mathcal{F}_t^Y] + \hat{g}_t \hat{h}_t \right) \right]$$

and since ξ_t belongs to a complete set, must have

$$f_t = \frac{\mathbb{E}[g(X_t)h(X_t)|\mathcal{F}_t^Y] - \hat{g}_t \hat{h}_t}{\gamma} = \frac{\widehat{gh}_t - \hat{g}_t \hat{h}_t}{\gamma}$$

which proves the theorem. \square

Existence of the solutions to the filtering SDE in theorem 5.2.1 is consequence of the fact that \hat{g}_t is on such solution. The uniqueness of solutions to the filtering SDE can be grouped in with proofs for uniqueness of Zakai equation (see Kurtz and Ocone [32] or Rozovsky [42]) because there is a one-to-one relationship between measure-valued solutions of the two SDEs.

5.2.1 Correlated Noise Filtering

Suppose that W_t and X_t are correlated so that for any function $g \in \mathcal{B}_b(Q)$ we have

$$\frac{1}{\gamma} \sum_n \Delta Y_{t_n} \Delta g(X_{t_n}) \xrightarrow{p} \langle W, g(X) \rangle_t \doteq \int_0^t \rho_s^g ds$$

where $0 = t_0 < t_1 < \dots < t_n = t$ and the limit is taken as $\sup_n (t_{n+1} - t_n) \rightarrow 0$. Then there is an added term in equation (5.3),

$$d\mathbb{E}[g(X_t)\xi_t] = \mathbb{E}[Qg(X_t)\xi_t] dt + \frac{i}{\gamma} \mathbb{E}[g(X_t)\xi_t \psi_t h(X_t)] dt + i \mathbb{E}[\xi_t \psi_t \rho_t^g] dt$$

and so for almost every $t \in [0, T]$ we have

$$\mathbb{E} \left[\xi_t \left(\gamma f_t - \mathbb{E}[g(X_t)h(X_t)|\mathcal{F}_t^Y] + \hat{g}_t \hat{h}_t - \gamma \hat{\rho}_t^g \right) \right] = 0$$

and so $f_t = \frac{1}{\gamma}(\widehat{gh}_t - \hat{g}_t \hat{h}_t + \gamma \hat{\rho}_t^g)$, and the nonlinear filter is

$$d\hat{g}_t = \widehat{Qg}_t dt + \frac{\widehat{gh}_t - \hat{g}_t \hat{h}_t + \gamma \hat{\rho}_t^g}{\gamma^2} d\nu_t.$$

5.2.2 The Kalman-Bucy Filter

Another big advantage to the innovations approach is in linear filtering. In particular, the case when the filtering problem consists of a system of linear SDEs. In this case, one needs to take some steps to verify that the filtering distribution is normal, but after doing so it is straight-forward to derive equations for the posterior's first and second moments.

Consider a non-degenerate linear observations model such that $h(x) = h \cdot x$ and $\gamma > 0$, with X_0 being Gaussian distributed, and with state-space generator

$$Q = \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} \cdot + ax \frac{\partial}{\partial x} \cdot$$

for all functions $g \in C^2(\mathbb{R})$, where σ and a are constant coefficients. The process

$$N_t \doteq \widehat{X}_t - \int_0^t \widehat{Q} \widehat{X}_s ds = \widehat{X}_t - a \int_0^t \widehat{X}_s ds$$

is square-integrable and a martingale, and if we extend proposition 5.1.2 for $h(x) = h \cdot x$ (see proposition 2.31 on page 34 of Bain and Crisan [5]), we can then write N_t using innovations Brownian motion,

$$N_t = \widehat{X}_0 + \frac{1}{\gamma} \int_0^t f_s d\nu_s$$

where $f_t \in L^2[0, T]$ and is \mathcal{F}_t^Y -adapted. Now, by simple stochastic calculus we can verify that X_t is given by

$$X_t = e^{at} X_0 + \sigma \int_0^t e^{a(t-s)} dB_s$$

where $B_t \perp W_t$ (i.e. the state-space noise is independent of the observation noise). Defining the estimation error ϵ_t as

$$\epsilon_t \doteq X_t - \widehat{X}_t.$$

and applying Itô's lemma, we have

$$d\epsilon_t = a\epsilon_t dt + \sigma dB_t - \frac{1}{\gamma} f_t d\nu_t,$$

which has a solution given by the integrating factor,

$$\epsilon_t = e^{at} \epsilon_0 + \sigma \int_0^t e^{a(t-s)} dB_s - \frac{1}{\gamma} \int_0^t e^{a(t-s)} f_s d\nu_s,$$

which is Gaussian distributed and uncorrelated with Y

$$\mathbb{E}[Y_s \epsilon_t] = \mathbb{E}[Y_s (X_t - \widehat{X}_t)] = \mathbb{E}[Y_s \mathbb{E}[X_t - \widehat{X}_t | \mathcal{F}_t^Y]] = 0$$

for all $s \leq t$. Now observe the following:

- (X_t, Y_s) are jointly Gaussian for any $s \leq t$,
- \hat{X}_t is Gaussian and a linear function of $\{Y_s\}_{s \leq t}$,
- therefore, (X_t, \hat{X}_t, Y_s) are jointly Gaussian for any $s \leq t$
- in particular (ϵ_t, Y_s) are jointly Gaussian and uncorrelated for any $s \leq t$.

Therefore, ϵ_t is independent of \mathcal{F}_t^Y (for a more detailed discussion see [5, 36]), and so the filter is Gaussian .

If we apply the filter in theorem 5.2.1 with $g(x) = x$, it yields $f_t = \frac{1}{\gamma} \mathbb{E}[\epsilon_t^2 | \mathcal{F}_t^Y] = \frac{1}{\gamma} \mathbb{E} \epsilon_t^2$, and an SDE for the evolution of the first filtering moment

$$d\hat{X}_t = a\hat{X}_t dt + \frac{h \cdot \mathbb{E} \epsilon_t^2}{\gamma^2} d\nu_t. \quad (5.4)$$

We can also apply Itô's lemma and then take expectations, which will result in a Riccati equation for the evolution of the filter's covariance,

$$\begin{aligned} d\mathbb{E}[\epsilon_t^2] &= 2a\mathbb{E}[\epsilon_t^2]dt + \sigma^2 dt + \mathbb{E}[f_t^2]dt \\ &= 2a\mathbb{E}[\epsilon_t^2]dt + \sigma^2 dt - \frac{h^2 \cdot \mathbb{E}[\epsilon_t^2]^2}{\gamma^2} dt \end{aligned} \quad (5.5)$$

where we have used the fact that $d\nu_t \cdot dB_t = 0$, and $\mathbb{E}[f_t^2] = \frac{1}{\gamma^2} \mathbb{E} [\mathbb{E}[\epsilon_t^2] \mathbb{E}[\epsilon_t^2]] = \frac{1}{\gamma^2} \mathbb{E}[\epsilon_t^2]^2$. Equations (5.4) and (5.5) are the Kalman filter.

Chapter 6

Numerical Methods for Approximating Nonlinear Filters

In practice, the exact filter can only be computed for models that are completely discrete, or for discrete-time linear Gaussian models in which the Kalman filter applies. In the other cases, the consistency of approximating schemes can be relatively trivial, while for others there is a fair amount of analysis required. In this lecture we present the general theory of Kushner [33] regarding the consistency of approximating filters, and also the Markov chain approximation methods of Dupuis and Kushner [21]. But first we present the following simple result regarding filter approximations:

Example 6.0.1. *Let X_t be an unobserved Markov process, and let the observation process Y_t be given by an SDE*

$$dY_t = h(t, X_t)dt + \gamma dW_t$$

where W_t is an independent Wiener process, and $\gamma > 0$. For any partition of the interval $[0, t]$ into N -many points, let $\mathcal{F}_t^N = \sigma\{Y_{t_n} : n \leq N\}$, and let the filtration generated by the continuum of observations be denoted by $\mathcal{F}_t^Y = \sigma\{Y_s : s \leq t\}$. Assuming that

$$\mathcal{F}_t^N \subset \mathcal{F}_t^{N+1} \subset \dots \subseteq \mathcal{F}_t^Y,$$

then from the continuity of Y_t we know that $\bigvee_{N=1}^{\infty} \mathcal{F}_t^N = \mathcal{F}_t^Y$. Then, by Lévy's 0-1 law we know that the conditional expectations converge,

$$\lim_N \mathbb{E}[g(X_t) | \mathcal{F}_t^N] = \mathbb{E} \left[g(X_t) \middle| \bigvee_{N=1}^{\infty} \mathcal{F}_t^N \right] = \mathbb{E}[g(X_t) | \mathcal{F}_t^Y].$$

for any integrable function $g(x)$. This clearly shows that the filter with discrete observations can be a consistent estimator of the filter with a continuum of observations. However, computing the filter with discrete observations may still require some approximations.

6.1 Approximation Theorem for Nonlinear Filters

In this section we present a proof of a theorem that essentially says: filtering expectations of bounded functions can be approximated by filters derived from models whose hidden state converges weakly to the true state. The theorem is presented in the context of continuous-time process with a continuum of observations, but can be reapplied in other cases with relatively minor changes.

For some $T < \infty$ and any $t \in [0, T]$, let X_t be an unobserved Markov process with generator Q with domain $\mathcal{B}(Q)$, and let $\mathcal{B}_b(Q)$ denote the subset of bounded functions in $\mathcal{B}(Q)$. For any function $g \in \mathcal{B}_b(Q)$, we have the following limit

$$\frac{\mathbb{E}[g(X_{t+\Delta t})|X_t = x] - g(x)}{\Delta t} \rightarrow Qg(x)$$

as $\Delta t \searrow 0$. The observed process Y_t is given by an SDE

$$dY_t = h(X_t)dt + \gamma dW_t \quad (6.1)$$

where W_t is an independent Wiener process, $\gamma > 0$ and h is a bounded function. For any time $t \in [0, T]$, let $\mathcal{F}_t^Y = \sigma\{Y_s : s \leq t\}$. From our study of Zakai equation we know that we can write the filtering expectation as

$$\mathbb{E}[g(X_t)|\mathcal{F}_t^Y] = \frac{\mathbb{E}[g(\tilde{X}_t)\tilde{M}_t|\mathcal{F}_t^Y]}{\mathbb{E}[\tilde{M}_t|\mathcal{F}_t^Y]} \quad (6.2)$$

where the paths of \tilde{X} have the same law as those of X but are independent of (X, Y) , and with \tilde{M}_t being the likelihood ratio

$$\tilde{M}_t \doteq \exp \left\{ \frac{1}{\gamma^2} \int_0^t h(\tilde{X}_s) dY_s - \frac{1}{2\gamma^2} \int_0^t h^2(\tilde{X}_s) ds \right\}.$$

The Zakai equation can provide us with an SDE for filtering expectations, but direct numerical quadrature methods to compute the Zakai equation may be difficult to justify.

6.1.1 Weak Convergence

In order to approximate the nonlinear filter, we will look to approximate X with a family of process $\{X^n\}$ which converge weakly to X . Let $\{\mathbb{P}^n\}$ be a family of measures on a metric space \mathcal{D} .

Definition 6.1.1. Weak Convergence. \mathbb{P}^n is said to converge weakly to \mathbb{P} if

$$\int f(x) d\mathbb{P}^n(x) \rightarrow \int f(x) d\mathbb{P}(x) \quad \text{as } n \rightarrow \infty$$

for any bounded continuous function $f : \mathcal{D} \rightarrow \mathbb{R}$. For the induced processes $\{X^n\}_n$, we denote weak convergence by writing $X^n \Rightarrow X$.

Definition 6.1.2. *Tightness.* We say that the family of measures $\{\mathbb{P}^n\}_n$ is tight if for any $\epsilon > 0$ there exists a compact set $K_\epsilon \subset \mathcal{D}$ such that

$$\inf_n \mathbb{P}^n(K_\epsilon) \geq 1 - \epsilon.$$

If so we also say that the induced processes $\{X^n\}_n$ are tight.

An important result regarding tightness is Prokhorov's theorem:

Theorem 6.1.1. *Prokhorov.* If \mathcal{D} is a complete and separable metric space, then the family $\{\mathbb{P}^n\}_n$ contained in the space of all probability measure on \mathcal{D} is relatively compact in the topology of weak convergence iff it is tight.

For the purposes of our study, we will consider processes X which are right continuous with left-hand limits (càdlàg). Let $D[0, T]$ denote the space of càdlàg functions from $[0, T]$ to \mathbb{R} , equipped with Skorohod topology. If a family of probability measures $\{\mathbb{P}^n\}_n$ is tight, then weak convergence of \mathbb{P}^n to the measure on X can be shown by verifying that the laws of the induced processes $\{X^n\}_n$ converge to the law of solutions to the associated martingale problem

$$\mathbb{E}^n \left[g(X_t^n) - g(X_s^n) - \int_s^t Qg(X_\tau^n) d\tau \middle| \mathcal{F}_s \right] \rightarrow 0$$

in probability as $n \rightarrow \infty$ for any $g(x) \in \mathcal{B}_b(Q)$, and for any $s \leq t$. It can be shown that the law of the solution to this martingale problem is unique.

Given a family of measure $\{\mathbb{P}^n\}_n$, an extremely useful tool is the Skorohod Representation Theorem, which says the following:

Theorem 6.1.2. *Skorohod Representation.* Let $\{\mathbb{P}^n\}_n$ be a family of measures on a complete and separable metric space. If \mathbb{P}^n converges weakly to \mathbb{P} , then there is a probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ with random variables \tilde{X}^n and \tilde{X} for which

- $\tilde{\mathbb{P}}(X^n \in A) = \mathbb{P}^n(A)$ for all n and any set $A \subset D[0, T]$,
- $\tilde{\mathbb{P}}(X \in A) = \mathbb{P}(A)$ for any set $A \subset D[0, T]$,
- and $\tilde{X}^n \rightarrow \tilde{X}$ $\tilde{\mathbb{P}}$ -a.s. as $n \rightarrow \infty$.

In particular, if $\tilde{X}^n \rightarrow \tilde{X} \in C[0, T]$ a.s. in the Skorohod topology, then the convergence holds uniformly in t . For a detailed treatment of weak convergence, the Skorohod topology, martingale problems, and Skorohod representations, see the book by Ethier and Kurtz [23] and the book by Yin and Zhang [47].

6.1.2 Consistency Theorem

Let $\{X^n\}_n$ be a family of random variables on the same probability space as (X, Y) , which also converge weakly to X . Let \tilde{X}^n denote a copy of X^n that is independent of (X, Y) , let the approximated likelihood be denoted by \tilde{M}_t^n ,

$$\tilde{M}_t^n \doteq \exp \left\{ \frac{1}{\gamma^2} \int_0^t h(\tilde{X}_s^n) dY_s - \frac{1}{2\gamma^2} \int_0^t h^2(\tilde{X}_s^n) ds \right\}$$

and define the approximated filtering expectation

$$\mathcal{E}_t^n [g(X_t)] \doteq \frac{\mathbb{E}[g(\tilde{X}_t^n) \tilde{M}_t^n | \mathcal{F}_t^Y]}{\mathbb{E}[\tilde{M}_t^n | \mathcal{F}_t^Y]} \quad \forall t \in [0, T] \quad (6.3)$$

for any function $g(x) \in \mathcal{B}_b(Q) \cap C(\mathcal{S})$. We then have the following theorem:

Theorem 6.1.3. *Given a family of process $\{X^n\}$ taking values in $D[0, T]$ which converge weakly in the Skorohod topology to $X \in C[0, T]$, the approximated filter in (6.3) will converge uniformly*

$$\sup_{t \leq T} |\mathcal{E}_t^n [g(X_t)] - \mathbb{E}[g(X_t) | \mathcal{F}_t^Y]| \rightarrow 0, \quad \text{for } X_t \in C[0, T]$$

in probability and in mean as $n \rightarrow \infty$, for any function $g(x) \in \mathcal{B}_b(Q) \cap C(\mathcal{S})$.

Proof. (taken from [33]) For the purposes of this proof, we can neglect the h^2 terms in M_t and \tilde{M}_t^n . Let ζ^1 and ζ^2 be bounded processes independent of W , and consider the following estimate,

$$V = \mathbb{E} \sup_{t \leq T} \left| \exp \left\{ \int_0^t \zeta_s^1 dW_s \right\} - \exp \left\{ \int_0^t \zeta_s^2 dW_s \right\} \right|. \quad (6.4)$$

We will use the following inequality that holds for real numbers a and b ,

$$|e^a - e^b| \leq |a - b|(e^a + e^b). \quad (6.5)$$

For any real-valued sub-martingale N_t , we have the following inequality (see [30]),

$$\mathbb{E} \sup_{t \leq T} N_t^2 \leq 4\mathbb{E} N_T^2. \quad (6.6)$$

Inequality (6.5) and a Schwarz inequality applied to (6.4) yields,

$$V^2 \leq \mathbb{E} \sup_{t \leq T} \left| \int_0^t (\zeta_s^1 - \zeta_s^2) dW_s \right|^2 \times \mathbb{E} \sup_{t \leq T} \left| \exp \left\{ \int_0^t \zeta_s^1 dW_s \right\} + \exp \left\{ \int_0^t \zeta_s^2 dW_s \right\} \right|^2 \quad (6.7)$$

By (6.6) the first term in (6.7) is bounded by $4\mathbb{E} \int_0^T |\zeta_s^1 - \zeta_s^2|^2 ds$. To bound the second term we use the fact that

$$\mathbb{E} \exp \left\{ \int_0^t \zeta_s^i dW_s \right\} = \mathbb{E} \exp \left\{ \frac{1}{2} \int_0^t |\zeta_s^i|^2 ds \right\} \quad \text{for } i = 1, 2$$

along with (6.6) and the fact that $\exp \left\{ \frac{1}{2} \int_0^t \zeta_s^i dW_s \right\}$ are bounded sub-martingales. Using these facts we can find a constant C that depends on T , ζ^1 and ζ^2 , such that

$$V^2 \leq C \cdot \mathbb{E} \int_0^T |\zeta_s^1 - \zeta_s^2|^2 ds. \quad (6.8)$$

To prove the theorem it suffices to show that

$$\sup_{t \leq T} \left| \mathbb{E}[g(\tilde{X}_t^n) \tilde{M}_t^n | \mathcal{F}_t^Y] - \mathbb{E}[g(\tilde{X}_t) \tilde{M}_t | \mathcal{F}_t^Y] \right| \rightarrow 0$$

in probability as $n \nearrow \infty$. From the boundedness of h and T , we know that

$$\mathbb{E} \sup_{t \leq T} (M_t^n)^2 + \mathbb{E} \sup_{t \leq T} (M_t)^2 < \infty. \quad (6.9)$$

Let \tilde{X} be a copy of X that is independent of (X, Y) . By the Skorokhod representation theorem we can assume W.L.O.G. that $\{\tilde{X}^n\}_n$ are defined on the same probability space as (\tilde{X}, X, Y) , that each \tilde{X}^n is independent of (X, Y) , and that $X^n \rightarrow X$ a.s. In fact, because we have assumed that X^n converges to a continuous function on the Skorokhod topology, we know that the convergence is uniform,

$$\sup_{t \leq T} |\tilde{X}_t^n - \tilde{X}_t| \rightarrow 0, \quad \text{a.s.}$$

as $n \rightarrow \infty$. In particular, $\sup_{t \leq T} |g(\tilde{X}_t^n) - g(\tilde{X}_t)| \tilde{M}_t^n \rightarrow 0$ a.s.

Taking expectations, we have

$$\begin{aligned} & \mathbb{E} \sup_{t \leq T} \left| \mathbb{E}[g(\tilde{X}_t^n) \tilde{M}_t^n | \mathcal{F}_t^Y] - \mathbb{E}[g(\tilde{X}_t) \tilde{M}_t | \mathcal{F}_t^Y] \right| \\ & \leq \mathbb{E} \sup_{t \leq T} \mathbb{E} \left[|g(\tilde{X}_t^n) - g(\tilde{X}_t)| \cdot \tilde{M}_t^n \middle| \mathcal{F}_t^Y \right] + \|g\|_\infty \mathbb{E} \sup_{t \leq T} \mathbb{E} \left[|\tilde{M}_t^n - \tilde{M}_t| \middle| \mathcal{F}_t^Y \right] \\ & \leq \mathbb{E} \left[\mathbb{E} \left[\sup_{t \leq T} \left(|g(\tilde{X}_t^n) - g(\tilde{X}_t)| \cdot \tilde{M}_t^n \right) \middle| \mathcal{F}_t^Y \right] \right] + \|g\|_\infty \mathbb{E} \left[\mathbb{E} \left[\sup_{t \leq T} |\tilde{M}_t^n - \tilde{M}_t| \middle| \mathcal{F}_t^Y \right] \right] \\ & \leq \mathbb{E} \left[\sup_{t \leq T} \left(|g(\tilde{X}_t^n) - g(\tilde{X}_t)| \cdot \tilde{M}_t^n \right) \right] + C \|g\|_\infty \mathbb{E} \left[\int_0^T |h(\tilde{X}_s^n) - h(\tilde{X}_s)|^2 ds \right]. \quad (6.10) \end{aligned}$$

The first term in (6.10) goes to zero as $n \rightarrow \infty$ after applying the bound in (6.9) and then by calling dominated convergence. The second term in (6.10) goes to zero because of the bound in (6.8). Therefore, the approximated filter converges in mean, and in probability as well. \square

Corollary 1. *To generalize theorem 6.1.3 for any $X \in D[0, T]$ such that $X^n \Rightarrow X$, we simply need to rework the end of the proof to show that the limit holds pointwise,*

$$|\mathcal{E}_t^n[g(X_t)] - \mathbb{E}[g(X_t)|\mathcal{F}_t^Y]| \rightarrow 0, \quad \text{for } X_t \in D[0, T]$$

in probability and in mean as $n \rightarrow \infty$, almost everywhere $t \in [0, T]$, and for any function $g(x) \in \mathcal{B}_b(Q) \cap C(\mathcal{S})$.

6.2 Markov Chain Approximations

Theorem 6.1.3 applies directly when observations are available as often as needed. In addition, a continuum of observations allows us a certain amount of flexibility in our choice of approximation scheme.

6.2.1 Approximation of Filters for Continuous-Time Markov Chains

Let X_t be a finite-state Markov chain. Consider a time step $\Delta t = \frac{1}{n}$. For n finite, take $k = 0, 1, 2, 3, 4, \dots, T/\Delta t - 1$ and denote the Markov chain ξ_k^n with transition probabilities

$$\mathbb{P}(\xi_{k+1}^n = i | \xi_k^n = j) = \left[e^{Q^* \Delta t} \right]_{ij}$$

for any $i, j \in \{1, \dots, m\}$. The Markov chain X_k^n is discrete but could be extended to $D[0, T]$ by taking $X_t^n = \sum_k \xi_k^n \mathbf{1}_{t \in [t_k, t_{k+1})}$, **but this is not a continuous-time Markov chain.** The non-Markov structure of X_t^n does not prevent us from applying theorem 6.1.3, but showing tightness and convergence of the martingale problem will be easier if we can find a Markovian approximation.

Let $k = 1, 2, 3, 4, \dots$, take $\nu_k \sim iid \exp(1)$, and set $\tau_k = \tau_{k-1} + \nu_k \Delta t$ with $\tau_0 = 0$. Now let ξ_k^n be the discrete Markov chain that we have already defined, but now consider all k up until $\tau_{k+1} \geq T$. A continuous-time approximation of X_t is then

$$X_t^n = \sum_{k: \tau_k < T} \mathbf{1}_{t \in [\tau_k, \tau_{k+1})} \xi_k^n.$$

To show that $\{X^n\}_n$ is compact in $D[0, T]$ we proceed as follows:

Take any $f \in D[0, T]$ and for any $i = 0, 1, 2, 3, 4, \dots$ let \mathcal{J}_i denote the time of the i th jump,

$$\mathcal{J}_{i+1}(f) = \inf\{t > \mathcal{J}_i(f) : f_t^n \neq f_{\mathcal{J}_i}^n\} \wedge T,$$

with $\mathcal{J}_0(f) = 0$. For general $f \in D[0, T]$ these \mathcal{J}_i 's may be infinitesimally small, but they will be informative for processes which approximated continuous-time Markov chains.

Next, for any $\delta > 0$ define the set

$$A_\delta = \{f \in \{1, \dots, m\} \text{ s.t. } |f_{\mathcal{J}_{i+1}} - f_{\mathcal{J}_i}| \geq \delta \quad \forall \mathcal{J}_{i+1} \leq T\}.$$

We can easily check that

$$\mathbb{P}(X^n \in A_\delta^c) \leq 1 - e^{\delta \max_j Q_{jj}} \leq -\delta \max_j Q_{jj} \ll 1$$

for any δ small enough. Furthermore, for any sequence $\{f^n\}_n \subset A_\delta$ we can find a subsequence such that for any i we have

$$\mathcal{J}_i(f^{n_\ell}) \rightarrow \mathcal{T}_i \in [i\delta, T]$$

and

$$f_{\mathcal{J}_i(f^{n_\ell})}^{n_\ell} \rightarrow a_i \in \{1, \dots, m\}$$

as $\ell \rightarrow \infty$. Therefore, $f_s^{n_\ell} \rightarrow a_i$ for $s \in [\mathcal{T}_i, \mathcal{T}_{i+1})$, which shows that there is subsequence that converges point-wise. Therefore, since $D[0, T]$ is equipped with the point-wise metric it follows that \bar{A}_δ is compact.

Now, if we look at the martingale problem, we have

$$\begin{aligned} \mathbb{E}g(X_t^n) &= \sum_{k:k\Delta t \leq t} \mathbb{E}\Delta g(X_{t_k}) + \mathbb{E}g(\bar{X}_0) \\ &= \sum_{k:k\Delta t \leq t} \mathbb{E}Qg(X_{t_k})\Delta t + \mathbb{E}g(X_0) + O(\Delta t) \end{aligned}$$

which converges to the martingale problem as $\Delta t \searrow 0$. Therefore, $X^n \Rightarrow X$ weakly in $D[0, T]$.

Then for any $s, t \in [0, T]$ with $s < t$, the change in the likelihood ratio for any path is given by

$$\begin{aligned} M_{s,t}^n = \exp \left\{ \frac{1}{\gamma^2} \sum_{k:s < \tau_k \leq t} h(X_{(\tau_{k-1} \vee s)}^n)(Y_{\tau_k} - Y_{(\tau_{k-1} \vee s)}) \right. \\ \left. - \frac{1}{2\gamma^2} \sum_{k:s < \tau_k \leq t} h^2(X_{\tau_{k-1} \vee s})(\tau_k - (\tau_{k-1} \vee s)) \right\}. \end{aligned}$$

At time t , let the approximating filtering mass function be denoted by $\omega_t = (\omega_t^1, \dots, \omega_t^m)$. Given ω_s we have the following recursion for the filtering mass:

$$\omega_t^i = \frac{\sum_j \mathbb{E} \left[\mathbf{1}_{\tilde{X}_t=i} \tilde{M}_{s,t}^n \middle| \mathcal{F}_t^Y \vee \{\tilde{X}_s = j\} \right] \omega_s^j}{\sum_j \mathbb{E} \left[\tilde{M}_{s,t}^n \middle| \mathcal{F}_t^Y \vee \{\tilde{X}_s = j\} \right] \omega_s^j}.$$

6.2.2 Approximation of Filter for System of SDEs

Let the filtering problem be as follows

$$\begin{aligned} dX_t &= a(X_t)dt + \sigma(x)dB_t \\ dY_t &= h(X_t)dt + \gamma dW_t \end{aligned}$$

with h bounded, $\gamma > 0$, and $W_t \perp B_t$. If $\sigma^2(x) \geq \frac{1}{n}|a(x)|$ for all $x \in \mathbb{R}$, we can approximate X with the a Markov chain X^n taking paths in $D[0, T]$ and X_t^n taking values in $\mathcal{S}^n = \{0, \pm\frac{1}{n}, \pm\frac{2}{n}, \dots\}$, defined as follows: $\Delta t^n(x) = \frac{1}{n^2\sigma^2(x)}$, and given $X_t^n = x$ the conditional probability distribution at time $t + \Delta t^n(x)$ is given by

$$\mathbb{P}\left(X_{t+\Delta t^n(x)}^n = x \pm \frac{1}{n} \middle| X_t^n = x\right) = \frac{\sigma^2(x) \pm \frac{1}{n}|a(x)|}{2\sigma^2(x)}.$$

We can construct a continuous-time Markov chain from this discrete-time Markov chain by using exponential arrivals as we did in with finite-state Markov chains. The subsequent process \bar{X}_t^n has the following differential,

$$\bar{X}_t^n = X_0 + \int_0^t a(\bar{X}_s^n)ds + \int_0^t \sigma(\bar{X}_s^n)d\omega_s + \epsilon^n(t)$$

where ω_t is an independent Wiener process and $\epsilon^n(t)$ is a semi-martingale such that

$$\mathbb{E} \sup_{t \leq T} (\epsilon^n(t))^2 \rightarrow 0$$

as $n \rightarrow \infty$, and by theorem 2.7b on page 27 of Ethier and Kurtz [23], $\{\bar{X}^n\}_n$ are tight for all $T < \infty$. This process is a locally consistent approximation to X on $[0, T]$.

6.2.3 Discrete-Time Obseavations

For general discrete observations models, theorem 6.1.3 applies for all $\{t_n\}_n \subset [0, T]$ for which each t_n is an observation time. For simplicity, suppose that Y_t is unobserved for $t \in (0, 1)$ and that the only observations are available at times $t = 0$ and $t = 1$,

$$Y_1 = Y_0 + \int_0^1 h(X_s)ds + W_1.$$

and assume that X_t is a Markov chain on a finite state-space. We write the filtering mass recursively as

$$\begin{aligned} \pi_1(x) &= \mathbb{P}(X_1 = x | Y_0, Y_1) \\ &= \frac{1}{c} \sum_{v \in \mathcal{S}} \mathbb{E} \left[\mathbb{P}(Y_1 | Y_0, \{\tilde{X}_s\}_{s \leq 1}) \middle| \mathcal{F}_1^Y \vee \{\tilde{X}_1 = x, \tilde{X}_0 = v\} \right] e^{Q^*1}(x|v) \pi_0(v) \end{aligned}$$

where c is a normalizing constant, $\mathcal{F}_1^Y = \sigma\{Y_0, Y_1\}$, the process \tilde{X} is a copy of X that is independent from (X, Y) , $e^{Q^*t}(\cdot | \cdot)$ is the transition kernel of X , and the likelihood function is

$$\mathbb{P}(Y_1|Y_0, \{x_s\}_{s \leq 1}) = \exp \left\{ -\frac{1}{2} \left(\frac{Y_1 - Y_0 - \int_0^1 h(x_s) ds}{\gamma} \right)^2 \right\}$$

for any path $\{x_s\}_{s \leq 1}$. We approximate X with a discrete-time Markov Chain X^n such that

$$\mathbb{P}(X_{k+1}^n = x | X_k^n = x') = e^{Q^*/n}(x|x')$$

and in this case, theorem 6.1.3 applies at time $t = 1$ because

$$\frac{1}{n} \sum_k h(X_k^n) \Rightarrow \int_0^1 h(X_s) ds$$

as $n \rightarrow \infty$.

In theory, the approximated filter will converge, however there are still computational issues because as n gets smaller we will need to devise a method to compute the expected likelihood

$$\mathbb{E} \left[\exp \left\{ -\frac{1}{2} \left(\frac{Y_1 - Y_0 - \frac{1}{n} \sum_{k=0}^{n-1} h(\tilde{X}_k^n)}{\gamma} \right)^2 \right\} \middle| \mathcal{F}_1^Y \vee \{\tilde{X}_n^n = x, \tilde{X}_0^n = v\} \right]$$

where \tilde{X}^n is a copy of X^n that is independent of (X, Y) . Monte Carlo methods can also be used, but the conditioning of the expectation on X_n^n might slow the convergence.

Chapter 7

Linear Filtering

Filtering in general linear models is perhaps the most widely applied branch of filtering, but in the context of linearity the term ‘filtering’ refers to something that is fundamentally different from the probabilistic models and equations that comprise what mathematicians refer to as ‘filtering theory.’ The methods are not Bayesian, and probability’s involvement can be minimal at times, but some of the most important ideas in filtering theory, such as the use of *innovations*, can be traced back to their pragmatic roots in signal processing and ‘linear’ filtering.

7.1 General Linear Filters

Let the integer $n \in \{0, 2, \dots, N - 1\}$ denote a time index. The simplest way to present a filtering problem is to identify a given measurement as a signal plus noise,

$$Y_n = X_n + W_n$$

where W is a noise component with positive covariance R ,

$$R_\ell \doteq \mathbb{E}W_{n+\ell}W_n$$

for some integer-valued lag ℓ . The noise can be considered idiosyncratic, essentially meaning that it is orthogonal to X ,

$$\mathbb{E}X_{n+\ell}W_n = 0$$

for any time n and any lag ℓ . For any linear filter $H : \mathbb{R}^N \rightarrow \mathbb{R}^M$, the impulse response is its convolution with the measurement

$$(H * Y)_k = (H * X)_k + (H * W)_k$$

where the convolution is a function of a shift k ,

$$(H * Y)_k = \sum_{n=0}^{M-1} X_n H_{n-k}.$$

If $k = 0$, the convolution can be thought of as the inner-product. For the filter H , the signal-to-noise ratio (SNR) is defined as ratio of the signal response over the noise response,

$$SNR_H \doteq \frac{\|H * X\|^2}{\mathbb{E}\|H * W\|^2}.$$

The goal of linear filtering is to estimate X with a projection of Y ,

$$\hat{X} \doteq H * Y,$$

or at least raise SNR so we are in a position better suited to make an estimate. We can consider such an estimate to be ‘optimal’ if we have chosen H for which SNR is maximized. If \hat{X} is an unbiased estimator, then the SNR of the optimal estimator will be greater than the SNR of the raw measurement,

$$SNR_Y = \frac{\|X\|^2}{\mathbb{E}\|W\|^2} < \frac{\|X\|^2}{\mathbb{E}\|\hat{X}_n - X_n\|^2} \doteq SNR_{\hat{X}}$$

and clearly, the major obstacles will be in finding the optimal linear filter. Obviously, if $\mathbb{E}X_n W_m = 0$ for all m, n , then it would make sense to take H to be some function that is known to have non-zero inner-product with X but is also known to be orthogonal to W . It might be difficult find (let alone to invert) such a filter. More importantly, one should notice that maximizing the $SNR_{\hat{X}}$ is the same as minimizing mean-square error (MSE),

$$MSE(\hat{X}) = \mathbb{E}\|X - \hat{X}\|^2 \approx \frac{1}{N} \sum_n |X_n - \hat{X}_n|^2.$$

7.1.1 Reed/Matched Filters

The Reed filter looks for the linear filter H that maps $X \mapsto \mathbb{R}$ with optimal SNR. Since R is positive-definite, there is an invertible matrix A such that

$$R = AA'$$

which we can use along with the Cauchy-Schwarz inequality to get the bound on SNR for the linear filter,

$$SNR_H = \frac{|H'X|^2}{H'RH} = \frac{|H'AA^{-1}X|^2}{H'AA'H} = \frac{|(A'H)'(A^{-1}X)|^2}{(A'H)'(A'H)}$$

$$\begin{aligned}
&\leq \frac{|(A'H)'(A'H)| \cdot |(A^{-1}X)'(A^{-1}X)|}{(A'H)'(A'H)} \\
&= X'(AA')^{-1}X = X'R^{-1}X
\end{aligned}$$

where H' , X' and A' are the transpose of there respective matrix/vector. The linear filter that achieves this upper bound is

$$H^{reed} = R^{-1}X,$$

yielding an optimal estimate as

$$\hat{X}^{reed} = \arg \max_x \left(\frac{x'H^{reed}}{(Y-x)'H^{reed}} \right),$$

but this will require a search over the signal domain. However, if we can parameterize the domain of the signal, it will be possible to compress our search into a simpler procedure that requires us to merely test the SNR of relatively few parameters. For instance, if we know a priori that X will have a significant response with a only a few of the Fourier basis functions, we can reduce an algorithm's search-time simply by searching over the domain of a few Fourier coefficients.

7.1.2 Fourier Transforms and Bandwidth Filters

Fourier transforms and fast-Fourier transform (FFT) algorithms can easily be used as linear filters. The Fourier basis functions are can be used for the spectral decomposition of periodic functions, but we can without loss of generality extend an observed finite vector (Y_0, \dots, Y_{N-1}) into a periodic function simply by concatenating a backwards copy. With periodicity in hand, we can use the FFT to filter-out frequencies which we have determined apriori to not be part of the signal. In other words, we apply an FFT to the observed data and then reconstruct the signal by only considering the inverse FFT of the coefficients that are within a bandwidth known apriori to be where the signal resides.

Let Y^* denote the Fourier transform of Y , defined as

$$Y_k^* = \sum_{n=0}^{N-1} e^{\frac{2\pi i}{N}kn} Y_n$$

and the inverse Fourier transform

$$Y_n = \frac{1}{N} \sum_{k=0}^{N-1} e^{-\frac{2\pi i}{N}kn} Y_k^*$$

which allows us to reconstruct the measurement. The central idea in a bandwidth filter is the notion that the signal lives in specific range of frequencies. For instance, from linearity we have

$$Y^* = X^* + W^*$$

and if we know that the support of X 's Fourier coefficients is contained in a set $K_0 \subset \{0, \dots, N-1\}$ such that

$$X_n = \frac{1}{N} \sum_{k \in K_0} e^{-\frac{2\pi i}{N} kn} X_k^*,$$

then we can construct an estimate based on the pertinent bandwidth(s).

$$\hat{X} = \frac{1}{N} \sum_{k \in K_0} e^{-\frac{2\pi i}{N} kn} Y_k^*.$$




Figure 7.1: The FFT of a random-walk considered to be the signal, and the FFT of an iid random variable that is considered to be the noise. The noise occupies the mid-level frequencies whereas the signal does not. Therefore, we construct a bandwidth filter simply by inverting the FFT without the mid-level coefficients.




Figure 7.2: The bandwidth filter where the mid-level frequencies are removed. The estimated signal is clearer than the raw measurement, and there is an increase in SNR, from $SNR_Y = \|X\|^2 / \frac{1}{N} \sum_n W_n^2 = 7.53$, to $SNR_{\hat{X}} = \|X\|^2 / \frac{1}{N} \sum_n (X_n - \hat{X}_n)^2 = 55.10$, and with $MSE = \frac{1}{N} \sum_n (X_n - \hat{X}_n)^2 = .0594$.

For instance, suppose X is a random-walk and Y equals X plus a considerable amount of noise,

$$\begin{aligned} X_n &= X_{n-1} + B_n \\ Y_n &= X_n + \gamma W_n \end{aligned}$$

where B_n and W_n are independent white noises, and $\gamma > 0$. We should try to identify the bandwidth(s) that contain the support of X 's Fourier coefficients simply by looking at the support of the FFT of a random-walk and comparing it to the FFT of noise. From figure 7.1 we see that the dominant Fourier coefficients of a random-walk are either in an extremely high or an extremely low bandwidth, whereas the Fourier coefficients of the noise are evenly distributed across all bandwidths. If we take K_0 to be the high and low frequencies that only contain noise, then the reconstructed signal will have a higher SNR,

$$\begin{aligned} SNR_Y &= \frac{\|X^*\|^2}{\mathbb{E}\|W^*\|^2} = \frac{\|X^*\|^2}{\mathbb{E}\|Y^* - X^*\|^2} = \frac{\|X^*\|^2}{\sum_{k=0}^{n-1} \mathbb{E}|Y_k^* - X_k^*|^2} \\ &= \frac{\|X^*\|^2}{\sum_{k \in K_0} \mathbb{E}|Y_k^* - X_k^*|^2 + \sum_{k \notin K_0} \mathbb{E}|Y_k^*|^2} < \frac{\|X^*\|^2}{\sum_{k \in K_0} \mathbb{E}|Y_k^* - X_k^*|^2} \\ &= \frac{\|X^*\|^2}{\sum_{k \in K_0} \mathbb{E}|\hat{X}_k^* - X_k^*|^2} = \frac{\|X^*\|^2}{\mathbb{E}\|\hat{X}^* - X^*\|^2} = SNR_{\hat{X}}. \end{aligned}$$

Indeed, as can be seen in figure 7.2, the signal becomes clearer as we eliminate the mid-level frequencies, and there is an increase in SNR from the 7.53 of the raw measurement, to 55.10 given by the bandwidth-filtered estimate, with a MSE=.0594.

7.1.3 Wavelet Filters

Wavelets are a tool that is useful in identifying the local behavior of a noise-corrupted signal. Measurements are often times contain adequate information for someone to decipher the underlying signal, usually because they can ignore noise and identify a movement in the measurement is caused by signal. This is precisely how a wavelet works: each wavelet represents a movement that the signal is capable of making, and any piece of the signal who's cross-product resonates with the wavelet is removed and placed in its respective spot as part of a noiseless reconstruction of the underlying signal.

A wavelet basis consists of a set of self-similar functions

$$\psi_n^{k\ell} = 2^{k/2} \psi_{2^k(n-\ell)}$$

for integers k and ℓ , where the unindexed function ψ is the 'mother-wavelet'. A useful wavelet has support that is small relative the length of the signal (e.g. $\text{supp}(\psi) \ll N$), and

by construction should sum to zero and have norm 1,

$$\sum_{n \in \text{supp}(\psi)} \psi_n = 0, \quad \sum_{n \in \text{supp}(\psi)} |\psi_n|^2 = 1,$$

with $\text{supp}(\psi)$ denoting the support of the wavelet. The wavelets are indexed by k and ℓ where k is a dilation and ℓ is a translation. The indices of the wavelets are chosen to form an orthonormal basis,

$$\sum_{n=0}^{N-1} \psi_n^{k\ell} \psi_n^{k'\ell'} = \mathbf{1}_{k=k'} \mathbf{1}_{\ell=\ell'},$$

and like any other spectral method we can reconstruct a function from its wavelet transform,

$$Y_n = \sum_{k,\ell} \langle Y, \psi^{k\ell} \rangle \psi_n^{k\ell}$$

for all $n \leq N - 1$, with $\langle \cdot, \cdot \rangle$ denoting inner-product. Not all wavelets can be used to form an orthonormal basis (e.g. the Mexican hat), but such wavelets should not be considered useless. Rather, a wavelet without an orthonormal basis simply requires that one use methods other than spectral decomposition.

For a signal processing problem, a particular wavelet is chosen for its generic resemblance to a local behavior of which the signal is capable. Essentially, we are taking a convolution of the function with wavelets of different thickness k , so that at each shift ℓ the local shape of the wavelet is given a chance to match itself to the input function. The wavelet-family that one uses will depend on the nature of the signal. Some possible wavelet to use in the construction of a discrete and orthogonal basis are symlets, coiflets, Daubechies, and Haar. The ‘mother wavelets’ for symlet24, a Daubechies24, a coiflet5, and a Haar, are shown in figure 7.3. Symlet, coiflet and Daubechie wavelets can be defined with their respective degree of differentiability. For instance, a family of symlet-5 wavelets are generated by a mother-wavelet that has at least 6 derivatives and so three first 5 wavelet moments are vanishing,

$$\sum_{n \in \text{supp}(\psi)} \psi_n n^p = 0$$

for $p = 0, 1, 2, 3, 4, 5$. In general, a wavelet that is $p + 1$ -times differentiability with fast enough decay in its tails will p -many vanishing moments.

When using wavelets to de-noise the random-walk example from section 7.1.2, there are numerous choices to make such as which wavelets to use and at what parameter values will we be fitting noise and not the signal. Figure 7.4 shows the wavelets’ ability to extract the signal from the noisy measurements, and table 7.1.3 shows how any of these wavelets does a better job de-noising the bandwidth filter from section 7.1.2.




Figure 7.3: Some examples of wavelets that can be used to construct an orthonormal basis.

Table 7.1: **Wavelet Denoising in Random-Walk Example.**

wavelet	SNR	MSE
coif2	64.90	0.0549
sym2	57.89	0.0581
db2	57.89	0.0581
coif5	59.32	0.0573
sym5	58.38	0.0578
db5	57.78	0.0581
haar	56.09	0.0589




Figure 7.4: The wavelets' de-noising of the measurement to uncover the signal is effective if we use the appropriate wavelet-family, and if we know what levels of the basis are not associated with the signal.

7.2 Linear Gaussian Models

A special case is when the impulse response is a linear model with Gaussian noise,

$$Y_n = X_n + W_n$$

where $R \doteq \mathbb{E}WW'$ is the variance/covariance matrix of a mean-zero Gaussian noise. We give ourselves a greater ability to infer the state of the signal simply by assuming that the noise is Gaussian. In the simplest case, just knowing the covariance properties of the system is enough to make a projection onto a basis of orthogonal basis, a projection that may even be a posterior expectation if the model can be shown to have a jointly-Gaussian structure. If we further assume that the signal evolves according to an independent Gaussian model we can apply a Kalman filter, which is extremely effective for tracking hidden Markov processes, particularly ones of multiple dimension.

7.2.1 The Wiener Filter

Given the data $Y = (Y_0, \dots, Y_N)'$, the signal $X = (X_0, \dots, X_N)'$ combines with a noise $W = (W_0, \dots, W_N)'$ so that

$$Y = X + W$$

where $\mathbb{E}W = 0$, the covariance matrix of the noise is

$$R \doteq \mathbb{E}WW',$$

and the covariance matrix of X is

$$Q \doteq \mathbb{E}[(X - \mathbb{E}X)X'].$$

The information introduced by Y can be encapsulated in **the innovation**,

$$V \doteq Y - \mathbb{E}X,$$

and the optimal linear estimate of X is its projection,

$$\mathcal{P}_Y X \doteq \mathbb{E}X + GV$$

where the matrix G is defined apriori in such a way as to make the projection error orthogonal to the posterior information:

$$\begin{aligned} 0 &= \mathbb{E}[(X - \mathcal{P}_Y X)Y'] = \mathbb{E}[(X - \mathbb{E}X - GV)Y'] \\ &= \mathbb{E}[(X - \mathbb{E}X)X'] + \mathbb{E}XW' - \mathbb{E}[G(Y - \mathbb{E}X)Y'] \end{aligned}$$

$$= Q - G(Q + R) \quad (*)$$

where we have assumed that $\mathbb{E}XW' = 0$ because noise by construction should be independent of the signal. If we solve $(*)$ we get

$$G = Q(Q + R)^{-1} \quad (7.1)$$

which is the optimal projection matrix. The Wiener filter is essentially a linear projection using the matrix in (7.1). Notice that we have made minimal assumptions about the distributions of the random variables; all we have assumed is that we know the mean and covariance structure of X and the driving noise in Y .

If we assume that (X, Y) are jointly Gaussian, then any random variable with the same distribution as X is equal in distribution to a random variable that is a linear sum of Y and another Gaussian component that is independent of Y ,

$$X = {}_d\mathbb{E}X + F_1(Y - \mathbb{E}X) + F_2Z$$

where F_1 and F_2 are non-random matrices of coefficients, and Z is mean-zero Gaussian and independent of Y . With this representation we have

$$\mathbb{E}[X|Y] = {}_d\mathbb{E}X + F_1V.$$

Now, by independence of Y and Z , we must have

$$0 = \mathbb{E}[(X - \mathbb{E}[X|Y])Y']$$

which leads to the solution $F_1 = G$ where G is the projection matrix given by (7.1). In general, the MSE is bounded below by that of the posterior mean,

$$\mathbb{E}\|X - \mathbb{E}[X|Y]\|^2 \leq \mathbb{E}\|X - \mathcal{P}_Y X\|^2.$$

But if Y and X are not jointly Gaussian, it can be shown that the MSE of the projection will be strictly greater than that of the posterior mean. In figure 7.5 the Wiener filter is used to track the random-walk example that was in section 7.1.2. For the random-walk example, the matrices are

$$Q = \begin{bmatrix} 1 & 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & 3 & \dots & 3 \\ 1 & 2 & 3 & 4 & \dots & 4 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & 4 & \dots & N \end{bmatrix}, \quad R = I_{N \times N}$$

which are ill-conditioned, but the round-off error is not significant for $N = 1000$. Indeed, the SNR and MSE of the Wiener filter is 85.79 and .0479, both are better than the best results among the wavelet and bandwidth filters (the best was the coiflet with vanishing moments which had $\text{SNR} = 64.90$ and $\text{MSE} = 0.0549$). This example illustrates how the Wiener filter is the optimal among all posterior estimators.




Figure 7.5: The Wiener filter applied to a random-walk. The $\text{SNR} = 85.69$ and $\text{MSE} = .0479$. The SNR/MSE is higher/lower than it was for the bandwidth filter since the Wiener filter is the optimal posterior estimator.

Remark 4. *The resemblance of the Wiener filter to a penalized and weighted least-squares problem is clear from the first order conditions of the following minimization,*

$$\min_X \left(\frac{1}{2} X' Q^{-1} X - X' (Q + R)^{-1} Y \right).$$

Remark 5. *When $X = (X_0, X_2, \dots, X_{N-1})$ is a realization from a a jointly Gaussian HMM, the Wiener filter returns not only the posterior mean, but also the path of X that is the maximum likelihood. In such cases, the estimator \hat{X}_n for $n < N - 1$ can be consider a smoothing rather than a filtering because it is an estimate of the state's past value,*

$$\hat{X}_n = \mathbb{E}[X_n | Y_{0:N-1}] \quad \text{for } n < N - 1.$$

Remark 6. *The numerical linear algebra for computing the Wiener filter requires no inversion of matrices, but only to solve two linear systems. Observe, \hat{X} is the solution to a linear system,*

$$(Q + R) \underbrace{Q^{-1} \hat{X}}_{=Z} = Y$$

so first we solve a linear system for Z

$$(Q + R)Z = Y$$

and then we solve for the filter,

$$\hat{X} = QZ.$$

7.2.2 The Kalman Filter

The Kalman filter can be thought of as a generalization of the Wiener filter, but for a model with a slightly more specific model for the signal. In fact, the Kalman filter is a filter for an HMM whose dynamics are Gaussian and fully linear,

$$\begin{aligned} X_n &= AX_{n-1} + B_n \\ Y_n &= HX_n + W_n \end{aligned}$$

where W and B are independent Gaussian random variables with covariance matrices

$$\begin{aligned} Q &= \mathbb{E}B_n B_n' \\ R &= \mathbb{E}W_n W_n' \end{aligned} \tag{7.2}$$

both of which are positive-definite, and the distribution of (X_0, Y_0) is a joint Gaussian. We can re-write this two equations as one linear system,

$$\begin{pmatrix} X_n \\ Y_n \end{pmatrix} = \begin{pmatrix} A & 0 \\ HA & 0 \end{pmatrix} \begin{pmatrix} X_{n-1} \\ Y_{n-1} \end{pmatrix} + \begin{pmatrix} \sqrt{Q} & 0 \\ H & \sqrt{R} \end{pmatrix} \begin{pmatrix} B_n \\ W_n \end{pmatrix} \tag{7.3}$$

which is clearly a non-degenerate Gaussian system. In fact (7.3) has a stationary mean if the number 1 is not included in the spectrum of A .

Given the data $Y_{0:n} = (Y_0, \dots, Y_n)$, the Kalman filter will find the optimal projection of X_n onto the Gaussian sub-space spanned by Y by iteratively refining the optimal projection of X_{n-1} onto the space spanned by $Y_{0:n-1}$. Furthermore, the optimal projection will be equivalent to the posterior mean because (X, Y) are jointly Gaussian.

In this case we can identify a sequence of Gaussian random variables that are **the innovations**

$$V_n \doteq Y_n - HAX_{n-1},$$

but the idea is essentially the same as it was in the Wiener filter. Initially, letting $\hat{X}_0 = \mathbb{E}[X_0|Y_0]$ we use a Wiener filter to get

$$\hat{X}_0 = \mathbb{E}X_0 + G_0(Y_0 - H\mathbb{E}X_0)$$

where $G_0 = \text{var}(X_0)(R + \text{var}(X_0))^{-1}$. Clearly, $X_0 - \hat{X}_0 \perp Y_0$, and we can easily check that (X_0, \hat{X}_0, Y_0) is jointly Gaussian, and so it follows that $X_0 - \hat{X}_0$ is independent of Y_0 , and so posterior covariance is not a random variable,

$$\Sigma_0 \doteq \mathbb{E} \left[(X_0 - \hat{X}_0)X_0' \middle| Y_0 \right] = \mathbb{E}(X_0 - \hat{X}_0)X_0'.$$

Now we proceed inductively to identify the filter of X_n given $Y_{0:n}$. Suppose we have obtained the filter up to time $n-1$ with posterior mean

$$\hat{X}_{n-1} \doteq \mathbb{E}[X_{n-1}|Y_{0:n-1}],$$

for which $X_{n-1} - \hat{X}_{n-1}$ is independent of $Y_{0:n-1}$, and with covariance matrix

$$\Sigma_{n-1} \doteq \mathbb{E} \left[(X_{n-1} - \hat{X}_{n-1})X_{n-1}' \middle| Y_{0:n-1} \right] = \mathbb{E} \left[(X_{n-1} - \hat{X}_{n-1})X_{n-1}' \right].$$

When the observation Y_n arrives, the optimal projection will be

$$\mathcal{P}_n X_n \doteq A\hat{X}_{n-1} + G_n V_n$$

where G_n is a projection matrix that is known at time $n-1$; it's known before Y_n has been observed. We can verify that

- the distribution of (X_n, Y_n) conditioned on $Y_{0:n-1}$ is jointly Gaussian, and
- that V_n is independent of $Y_{0:n-1}$,

and therefore it follows that $\hat{X}_n = \mathcal{P}_n X_n$. We can also write the following expression for the prediction covariance matrix,

$$\begin{aligned} \Sigma_{n|n-1} &\doteq \mathbb{E} \left[(X_n - A\hat{X}_{n-1})X_n' \middle| Y_{0:n-1} \right] \\ &= \mathbb{E} \left[(AX_{n-1} + B_n - A\hat{X}_{n-1})(AX_{n-1} + B_n)' \right] \\ &= A\Sigma_{n-1}A' + Q \end{aligned}$$

and from the orthogonality of the projection residual to the data, we should have a projection matrix that satisfies the following equation,

$$0 = \mathbb{E} \left[(X_n - \hat{X}_n)Y_n' \middle| Y_{0:n-1} \right]$$

$$\begin{aligned}
&= \mathbb{E} \left[(X_n - A\hat{X}_{n-1})X_n' H' \middle| Y_{0:n-1} \right] - \mathbb{E} \left[G_n \left(Y_n - HA\hat{X}_{n-1} \right) Y_n' \middle| Y_{0:n-1} \right] \\
&= \Sigma_{n|n-1} H' - G_n (H \Sigma_{n|n-1} H' + R).
\end{aligned}$$

We solve this equation to obtain the optimal projection matrix, also known as **the Kalman filter Gain matrix**

$$G_n = \Sigma_{n|n-1} H' (H \Sigma_{n|n-1} H' + R)^{-1} \quad (7.4)$$

and using the gain matrix we can write the posterior mean as a recursive function of the innovation and the previous time's posterior mean

$$\hat{X}_n = A\hat{X}_{n-1} + G_n V_n. \quad (7.5)$$

Furthermore, we can verify that the conditional distribution of (X_n, \hat{X}_n, Y_n) is jointly Gaussian, and since $\mathbb{E}[(X_n - \hat{X}_n)Y_m'] = 0$ for all $m \leq n$, it follows that $X_n - \hat{X}_n$ is independent of $Y_{0:n}$. Therefore, the covariance matrix is not a function of the data

$$\begin{aligned}
\Sigma_n &= \mathbb{E} \left[(X_n - \hat{X}_n)X_n' \right] = \mathbb{E} \left[(X_n - \hat{X}_n)X_n' \middle| Y_{0:n-1} \right] \\
&= \mathbb{E} \left[(X_n - A\hat{X}_{n-1})X_n' \middle| Y_{0:n-1} \right] - G_n \mathbb{E} \left[(Y_n - HA\hat{X}_{n-1})X_n' \middle| Y_{0:n-1} \right] \\
&= A\Sigma_{n-1}A' + Q - G_n H (A\Sigma_{n-1}A' + Q) = (I - G_n H)\Sigma_{n|n-1}.
\end{aligned}$$

To summarize, we have shown that $X_n - \hat{X}_n \sim N(0, \Sigma_n)$ and independent of $Y_{0:n}$, and from equations (7.4), (7.5) along with the equations for $\Sigma_{n|n-1}$ and Σ_n we have **the Kalman filter** at time n

$$\begin{aligned}
\Sigma_{n|n-1} &= A\Sigma_{n-1}A' + Q \\
G_n &= \Sigma_{n|n-1} H' (H \Sigma_{n|n-1} H' + R)^{-1} \\
\hat{X}_n &= A\hat{X}_{n-1} + G_n V_n \\
\Sigma_n &= (I - G_n H)\Sigma_{n|n-1}
\end{aligned}$$

so that the posterior density of X_n is

$$p(X_n \in dx | Y_{0:n}) = \frac{1}{(2\pi|\Sigma_n|)^{d/2}} \exp \left\{ -\frac{1}{2}(x - \hat{X}_n)' \Sigma_n^{-1} (x - \hat{X}_n) \right\}$$

where d is the dimension such that $X_n \in \mathbb{R}^d$.

In figure 7.6 we see the Kalman filter's ability to track the same random-walk example on which we test the filters from sections 7.1.2, 7.1.3 and 7.2.1.




Figure 7.6: The Kalman filter applied to the random-walk example. The path of Kalman filter estimates is not optimal as a whole, but each \hat{X}_n returned by the Kalman filter is optimal given the information $Y_{0:n}$.

The random-walk model is

$$\begin{aligned} X_n &= X_{n-1} + B_n \\ Y_n &= X_n + \gamma W_n. \end{aligned}$$

and Kalman filter for the random-walk is

$$\begin{aligned} G_n &= (\Sigma_{n-1} + 1) (\Sigma_{n-1} + 1 + \gamma^2)^{-1} \\ \hat{X}_n &= \hat{X}_{n-1} + G_n(Y_n - \hat{X}_{n-1}) \\ \Sigma_n &= (1 - G_n) (\Sigma_{n-1} + 1). \end{aligned}$$

Given $Y_{0:N}$, the Kalman filter returns an optimal estimator of X_N , not the entire path taken by X . Indeed, the Kalman filter's path has $\text{SNR} = 43.88$ and $\text{MSE} = 0.0668$, neither of which are better than the other filters. But this example should not be evidence for a dismissal of the Kalman filter, it simply shows that ex-ante estimation of the entire path of the signal is not its specialty.

The Kalman filter is far superior to the other filters we've discussed when it is applied to problems where X_n is a multidimensional vector. When each observation is a vector, the curse of dimensionality makes it impossible to work with the basis' required for bandwidth and wavelets, and the size of the matrices needed for the Wiener filter also be prohibitively large. On the other hand, the Kalman filter works efficiently and in real-time.

Remark 7. *For HMMs, the Kalman filter and the Wiener filter coincide in their estimates of the latest value of the signal,*

$$\hat{X}_{N-1}^{\text{wiener}} = \mathbb{E}[X_{N-1} | Y_{0:N-1}] = \hat{X}_{N-1}^{\text{kalman}}.$$

Remark 8. *The Kalman filter is indeed capable of handling signals of high dimension, but there does not exist a general procedure for avoiding the explicit computation of the matrix inverse when computing the gain matrix. Sometimes this inverse may be manageable, but limitations in our ability to compute matrix inverse represent the upper-bound on the Kalman filter's capacity.*

Chapter 8

The Baum-Welch & Viterbi Algorithms

Filtering equations for the class of fully-discrete HMMs are relatively simple to derive through Bayesian manipulation of the posteriors. These discrete algorithms are interesting because they embody the most powerful elements of HMM theory in a very simple framework. The methods are readily-implementable and have become the workhorse in applied areas where machine learning algorithms are needed. The algorithms for filtering, smoothing and parameter estimation are analogous to their counterparts in continuous models, but the theoretical background required for understanding is minimal in the discrete setting.

8.1 Equations for Filtering, Smoothing & Prediction

Let n denote a discrete time, and suppose that X_n is an unobserved Markov chain taking values in a discrete state-space denoted by \mathcal{S} . Let Λ denote X_n 's kernel of transition probabilities so that

$$\mathbb{P}(X_{n+1} = x) = \sum_{v \in \mathcal{S}} \Lambda(x|v) \mathbb{P}(X_n = v)$$

for any $x \in \mathcal{S}$, and $\mathbb{P}(X_0 = x) = p_0(x)$.

Noisy measurements are taken in the form of a process Y_n which is a nonlinear function of X_n , plus some noise,

$$Y_n = h(X_n) + W_n$$

where W_n is an iid Gaussian random variable with mean zero and variance $\gamma^2 > 0$. The main feature of this discrete model is the **memoryless-channel** which allows the process to ‘forget the past’:

$$\mathbb{P}(Y_n, X_n = x | X_{n-1} = v, Y_{0:n-1}) = \mathbb{P}(Y_n | X_n = x) \Lambda(x|v)$$

for any $n \geq 0$ and for all $x, v \in \mathcal{S}$.

8.1.1 Filtering

The filtering mass function is

$$\pi_n(x) \doteq \mathbb{P}(X_n = x | Y_{0:n})$$

for all $x \in \mathcal{S}$. Through an application of Bayes rule along with the properties of the HMM, we are able to break down π_n as follows,

$$\begin{aligned} \pi_n(x) &= \frac{\mathbb{P}(X_n = x, Y_{0:n})}{\mathbb{P}(Y_{0:n})} \\ &= \frac{\sum_{v \in \mathcal{S}} \mathbb{P}(Y_n, X_n = x | X_{n-1} = v, Y_{0:n-1}) \mathbb{P}(X_{n-1} = v, Y_{0:n-1})}{\mathbb{P}(Y_{0:n})} \\ &= \frac{\mathbb{P}(Y_n | X_n = x) \sum_{v \in \mathcal{S}} \mathbb{P}(X_n = x | X_{n-1} = v) \mathbb{P}(X_{n-1} = v, Y_{0:n-1})}{\mathbb{P}(Y_{0:n})} \\ &= \frac{\mathbb{P}(Y_n | X_n = x) \sum_{v \in \mathcal{S}} \Lambda(x|v) \mathbb{P}(X_{n-1} = v | Y_{0:n-1})}{\mathbb{P}(Y_n | Y_{0:n-1})} \\ &= \frac{\mathbb{P}(Y_n | X_n = x) \sum_{v \in \mathcal{S}} \Lambda(x|v) \pi_{n-1}(v)}{\sum_{x \in \mathcal{S}} \text{numerator}} \end{aligned}$$

where the memoryless-channel allows for the conditioning that occurs between the second and third lines. This recursive breakdown of the filtering mass is the **forward Baum-Welch Equation**, and can be written explicitly for the the system with Gaussian observation noise

$$\pi_n(x) = \frac{1}{c_n} \psi_n(x) \sum_{v \in \mathcal{S}} \Lambda(x|v) \pi_{n-1}(v) \quad (8.1)$$

where c_n is a normalizing constant, and ψ_n is a likelihood function

$$\psi_n(x) \doteq \mathbb{P}(Y_n | X_n = x) = \exp \left\{ -\frac{1}{2} \left(\frac{Y_n - h(x)}{\gamma} \right)^2 \right\}.$$

Equation (8.1) is convenient because it keeps the distribution updated without having to recompute old statistics as new data arrives. In ‘real-time’ it is efficient to use this algorithm to keep track of X ’s latest movements, but older filtering estimates will not be optimal after new data has arrived. The smoothing distribution must be used to find the optimal estimate of X at some time in the past.

8.1.2 Smoothing

For some time $N > n$ up to which data has been collected, the smoothing mass function is

$$\pi_{n|N}(x) \doteq \mathbb{P}(X_n = x | Y_{0:N}).$$

Through an application of Bayes rule along with the properties of the model, the smoothing mass can be written as follows,

$$\begin{aligned} \pi_{n|N}(x) &= \frac{\mathbb{P}(Y_{n+1:N} | X_n = x) \pi_n(x)}{\mathbb{P}(Y_{n+1:N} | Y_{0:n})} \\ &= \frac{\sum_{v \in \mathcal{S}} \mathbb{P}(Y_{n+1:N} | X_{n+1} = v, X_n = x) \Lambda(v|x) \pi_n(x)}{\mathbb{P}(Y_{n+1:N} | Y_{0:n})} \\ &= \frac{\sum_{v \in \mathcal{S}} \mathbb{P}(Y_{n+2:N} | X_{n+1} = v) \psi_{n+1}(v) \Lambda(v|x) \pi_n(x)}{\mathbb{P}(Y_{n+2:N} | Y_{0:n+1}) \mathbb{P}(Y_{n+1} | Y_{0:n})} \\ &= \frac{\sum_{v \in \mathcal{S}} \mathbb{P}(Y_{n+2:N} | X_{n+1} = v) \psi_{n+1}(v) \Lambda(v|x) \pi_n(x)}{\mathbb{P}(Y_{n+2:N} | Y_{0:n+1}) c_{n+1}}. \quad (*) \end{aligned}$$

where c_{n+1} is the normalizing constant from equation (8.1). Now suppose that we define a likelihood function for the events after time n ,

$$\alpha_n^N(x) = \frac{\mathbb{P}(Y_{n+1:N} | X_n = x)}{\mathbb{P}(Y_{n+1:N} | Y_{0:n})}$$

for $n < N$ with the convention that $\alpha_N^N \equiv 1$. Then the smoothing mass can be written as the product of the filtering mass with α

$$\pi_{n|N}(x) = \alpha_n^N(x) \pi_n(x)$$

and from (*) we can see that α_n^N is given recursively by a **backward Baum-Welch Equation**

$$\alpha_n^N(x) = \frac{1}{c_{n+1}} \sum_{v \in \mathcal{S}} \alpha_{n+1}^N(v) \psi_{n+1}(v) \Lambda(v|x). \quad (8.2)$$

Clearly, computation of the smoothing distribution requires a computation of all filtering distribution up to time N followed by the backward recursion to compute α^N . In exchange for doing this extra work, the sequence of X 's estimates will suggest a path taken by X that is more plausible than the path suggested by the filtering estimates.

8.1.3 Prediction

The prediction distribution is easier to compute than smoothing. For $n < N$, the prediction distribution is

$$\pi_{N|n}(x) \doteq \mathbb{P}(X_N = x | Y_{0:n})$$

and is merely computed by extrapolating the filtering distribution,

$$\pi_{N|n}(x) = \sum_{v \in \mathcal{S}} \Lambda(x|v) \pi_{N-1|n}(v) = \sum_{v \in \mathcal{S}} \Lambda^{N-n}(x|v) \pi_n(v)$$

where Λ^{N-n} denotes the transition probability over $N - n$ time steps.

If X_n is a positive recurrent Markov chain, then there is an invariant and the prediction distribution will converge to as $N \rightarrow \infty$. In some cases, the rate at which this convergence occurs will be proportional to the spectral gap in Λ .

Suppose X_n can take one of m -many finite-state, and is a recurrent Markov chain with only 1 communication class. Let $\Lambda \in \mathbb{R}^{m \times m}$ be the matrix of transition probabilities for X , and suppose that $\Lambda_{ji} > 0$ so that

$$\mathbb{P}(X_{n+1} = x_i | X_n = x_j) = \Lambda_{ji} > 0$$

for all $i, j \leq m$. Then the prediction distribution is

$$\pi_{N|n} = \pi_n \Lambda^{N-n}$$

and will converge exponentially fast to the invariant measure with a rate proportional to the second eigenvalue of Λ . To see why this is true, consider the basis of eigenvectors $(\mu_i)_{i \leq m}$ of Λ , some of which may be generalized,

$$\mu_{i+1}(\Lambda - \beta_i I) = \mu_i$$

for some $i \geq 1$. Assuming that μ_1 is the unique invariant mass function of X_n , we have $\mu_1 \Lambda = \mu_1$. By the Perron-Frobenius Theorem we can sort the eigenvalues so that $1 = \beta_1 > |\beta_2| \geq |\beta_3| \geq \dots \geq |\beta_m|$, and we know that β_1 is a simple root of the characteristic polynomial and therefore μ_1 is not a generalized eigenvector. From here we can see that

$$\begin{aligned} -\frac{1}{k} \log \|\pi_n \Lambda^k - \mu_1\| &= -\frac{1}{k} \log \|(\pi_n - \mu_1) \Lambda^k\| = -\frac{1}{k} \log \|(a_1 \mu_1 + a_2 \mu_2 + \dots a_m \mu_m) \Lambda^k\| \\ &= -\frac{1}{k} \log \|a_1 \mu_1 + a_2 \beta_2^k \mu_2 + \dots a_m \mu_m \Lambda^k\| \sim \frac{1}{k} \log \left(1 + a'_2 |\beta_2^k| \right) \sim |\beta_2| \end{aligned}$$

as $k \rightarrow \infty$. The spectral gap of Λ is $1 - |\beta_2|$, and from the convergence rate we see that a greater spectral gap means that the prediction distribution will take less time to converge to the invariant measure. In general, the Perron-Frobenius theorem can be applied to a recurrent finite-state Markov chain provided that there is some integer $k < \infty$ for which $\Lambda_{ji}^k > 0$ for all $i, j \leq m$.

8.2 Baum-Welch Algorithm for Learning Parameters

It is not very realistic to assume that we have apriori knowledge of the HMM that is completely accurate. However, stationarity of X means that we are observed repeated behavior of X , albeit through noisy measurements, but nevertheless we should be able to judge the frequencies with which X occupies parts of the state-space and the frequencies with which it moves about.

If we have already computed the smoothing distribution based on a model that is ‘close’ in some sense, then we should have

$$\frac{1}{N} \sum_{n=1}^N \mathbb{P}(X_n = x | Y_{0:N}) \approx \mu(x) \quad (8.3)$$

$$\frac{1}{N} \sum_{n=1}^N \mathbb{P}(X_n = x, X_{n-1} = v | Y_{0:N}) \approx \Lambda(x|v)\mu(v) \quad (8.4)$$

where μ is the stationary law of X . With the Baum-Welch algorithm, we can in fact employ some optimization techniques to find a sequence of model estimates which are of increasing likelihood, and it turns out that the (8.3) and (8.4) are similar to the optimal improvement in selecting the sequence of models.

Consider two model parameters θ and θ' . The Baum-Welch algorithm uses the Kullback-Leibler divergence to compare the two models,

$$\begin{aligned} 0 \leq D(\theta || \theta') &= \sum_{\vec{x} \in \mathcal{S}^{N+1}} \frac{\mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N})}{\mathbb{P}^\theta(Y_{0:N})} \log \left(\frac{\mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N}) \mathbb{P}^{\theta'}(Y_{0:N})}{\mathbb{P}^{\theta'}(X_{0:N} = \vec{x}, Y_{0:N}) \mathbb{P}^\theta(Y_{0:N})} \right) \\ &= \log \left(\frac{\mathbb{P}^{\theta'}(Y_{0:N})}{\mathbb{P}^\theta(Y_{0:N})} \right) + \sum_{\vec{x} \in \mathcal{S}^{N+1}} \frac{\mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N})}{\mathbb{P}^\theta(Y_{0:N})} \log \left(\frac{\mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N})}{\mathbb{P}^{\theta'}(X_{0:N} = \vec{x}, Y_{0:N})} \right). \end{aligned}$$

If we set

$$Q(\theta || \theta') \doteq \sum_{\vec{x} \in \mathcal{S}^{N+1}} \mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N}) \log \left(\mathbb{P}^{\theta'}(X_{0:N} = \vec{x}, Y_{0:N}) \right),$$

we then have a simplified expression,

$$0 \leq D(\theta || \theta') = \log \left(\frac{\mathbb{P}^{\theta'}(Y_{0:N})}{\mathbb{P}^\theta(Y_{0:N})} \right) + \frac{Q(\theta || \theta) - Q(\theta || \theta')}{\mathbb{P}^\theta(Y_{0:N})}$$

and rearranging the inequality we have

$$\frac{Q(\theta || \theta') - Q(\theta || \theta)}{\mathbb{P}^\theta(Y_{0:N})} \leq \log \left(\frac{\mathbb{P}^{\theta'}(Y_{0:N})}{\mathbb{P}^\theta(Y_{0:N})} \right),$$

from which we see that $Q(\theta||\theta') > Q(\theta||\theta)$ implies that θ' has greater likelihood than θ . The Baum-Welch algorithm uses this inequality as the basis for a criteria to iteratively refine the estimated model parameter. The algorithm obtains a sequence $\{\theta^\ell\}_\ell$ for which $Q(\theta^{\ell-1}||\theta^\ell) \geq 0$, and so their likelihoods are increasing but bounded,

$$\mathbb{P}^{\theta^{\ell-1}}(Y_{0:N}) \leq \mathbb{P}^{\theta^\ell}(Y_{0:N}) \leq \mathbb{P}^{\hat{\theta}^{mle}}(Y_{0:N}),$$

where $\hat{\theta}^{mle}$ is the maximum likelihood estimate of θ . Therefore, $\{\theta^\ell\}_\ell$ will have a limit at θ^* such that

$$\mathbb{P}^{\theta^*}(Y_{0:N}) = \lim_{\ell} \mathbb{P}^{\theta^\ell}(Y_{0:N}),$$

but it may be the case that $\mathbb{P}^{\theta^*}(Y_{0:N}) < \mathbb{P}^{\hat{\theta}^{mle}}(Y_{0:N})$ (see figure 8.1).




Figure 8.1: A sequence of Baum-Welch parameter estimates with increasing likelihood, but the sequence is caught at a local maximum.

In doing computations, a maximum (perhaps only a local maximum) of $Q(\theta|| \cdot)$ needs to be found. First-order conditions are good technique for finding one, and using the HMM we can expand $Q(\theta||\theta')$ into an explicit form,

$$Q(\theta||\theta') = \sum_{\vec{x} \in \mathcal{S}^{N+1}} \mathbb{P}^\theta(X_{0:N} = \vec{x}, Y_{0:N}) \left\{ \psi_0^{\theta'}(\vec{x}_0) p_0^{\theta'}(\vec{x}_0) + \sum_{n=1}^N \log \left(\psi_n^{\theta'}(\vec{x}_n) \Lambda^{\theta'}(\vec{x}_n | \vec{x}_{n-1}) \right) \right\}, \quad (8.5)$$

from which we see that it is possible to differentiate with respect to θ' , add the Lagrangians, and then solve for the optimal model estimate.

The Baum-Welch algorithm is equivalent to the expectation-maximization (EM) algorithm; the EM algorithm maximizes the expectation of the log-likelihood function which is equivalent to maximizing Q ,

$$\theta^\ell = \arg \max_{\theta} \mathbb{E}^{\theta^{\ell-1}} \left[\log \left(\mathbb{P}^{\theta}(Y_{0:N}, X_{0:N}) \right) \middle| Y_{0:N} \right] = \arg \max_{\theta} Q(\theta^{\ell-1} \| \theta).$$

8.2.1 Model Re-Estimation for Parametric Transition Probabilities

Suppose that $X_n \in \mathbb{Z}$, with transition probabilities parameterized by $\theta \in (0, \infty)$ so that

$$\mathbb{P}(X_{n+1} = i | X_n = j) = \frac{1}{c(\theta)} \exp\{-\theta|i-j|^2\}, \quad \forall i, j \in \mathbb{Z},$$

where $c(\theta) = \sum_{i=-\infty}^{\infty} \exp\{-\theta|i-j|^2\}$. Ignoring the parts that do not depend on θ' , the log-likelihood is

$$Q(\theta \| \theta') = - \sum_{n=1}^N \mathbb{E}^{\theta} \left[\theta' |X_n - X_{n-1}|^2 + \log c(\theta') \middle| \mathcal{F}_N^Y \right],$$

and if we differentiate with respect to θ' we have the following first-order conditions,

$$\frac{\partial}{\partial \theta'} Q(\theta \| \theta') = - \sum_{n=1}^N \mathbb{E}^{\theta} \left[|X_n - X_{n-1}|^2 - \frac{\sum_i |i-j|^2 \exp\{-\theta|i-j|^2\}}{c(\theta')} \middle| \mathcal{F}_N^Y \right] = 0$$

for any $j \in \mathbb{Z}$. The solution to the first-order conditions is θ' that satisfies

$$\mathbb{E}^{\theta'} \left[|X_1 - X_0|^2 \middle| X_0 = j \right] = \frac{1}{N} \sum_{n=1}^N \mathbb{E}^{\theta} \left[|X_n - X_{n-1}|^2 \middle| \mathcal{F}_N^Y \right]$$

for any j .

8.2.2 Model Re-Estimation for Finite-State Markov Chains

Suppose $X_n \in \mathcal{S} = \{1, \dots, m\}$, so that

$$\mathbb{P}(X_{n+1} = i | X_n = j) = \Lambda_{ji}$$

for all $i, j \in \mathcal{S}$. We will look for a sequence $\Lambda^{(\ell)}$ which maximizes $Q(\Lambda^{(\ell-1)} \| \cdot)$ subject to the constraints $\sum_i \Lambda_{ji} = 1$ for all $j \leq m$. Letting δ_j be the Lagrange multiplier for the j th constraint, the first order conditions are then,

$$\frac{\partial}{\partial \Lambda_{ji}} \left(Q(\Lambda^{(\ell-1)} \| \Lambda) - \delta_j \sum_r \Lambda_{jr} \right) = \frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda) - \delta_j = 0. \quad (**)$$

Multiplying by Λ_{ji} and summing over i the expression in (**) becomes

$$0 = \sum_i \Lambda_{ji} \left(\frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda) - \delta_j \right) = \sum_i \Lambda_{ji} \frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda) - \delta_j$$

which means $\delta_j = \sum_i \Lambda_{ji} \frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda)$. By multiplying (**) by Λ_{ji} and then rearranging terms it is found that the optimal $\Lambda_{ji}^{(\ell)}$ must be chosen among the set of Λ 's such that

$$\Lambda_{ji} = \frac{\Lambda_{ji} \frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda)}{\sum_r \Lambda_{jr} \frac{\partial}{\partial \Lambda_{jr}} Q(\Lambda^{(\ell-1)} \| \Lambda)}. \quad (8.6)$$

Now, using the expansion in (8.5), the derivative of $Q(\Lambda^{(\ell-1)} \| \Lambda)$ with respect to Λ_{ji} can be computed as follows:

$$\begin{aligned} \frac{\partial}{\partial \Lambda_{ji}} Q(\Lambda^{(\ell-1)} \| \Lambda) &= \mathbb{E} \left[\frac{\partial}{\partial \Lambda_{ji}} \log \mathbb{P}(Y_{0:N}, X_{0:N} | \Lambda) \middle| Y_{0:N}, \Lambda^{(\ell-1)} \right] \\ &= \mathbb{E} \left[\sum_{n=1}^N \frac{1}{\Lambda_{ji}} \mathbf{1}_{\{X_n=i, X_{n-1}=j\}} \middle| Y_{0:N}, \Lambda^{(\ell-1)} \right] = \frac{1}{\Lambda_{ji}} \sum_{k=1}^N \mathbb{P}(X_n = i, X_{n-1} = j | Y_{0:N}, \Lambda^{(\ell-1)}) \end{aligned}$$

and by plugging this into equation (8.6) it is easily seen that the solution is

$$\Lambda_{ji}^{(\ell)} = \frac{\sum_{n=1}^N \mathbb{P}(X_n = i, X_{n-1} = j | Y_{0:N}, \Lambda^{(\ell-1)})}{\sum_i \text{numerator}} \quad (8.7)$$

where $\mathbb{P}(X_n = i, x_{n-1} = j | Y_{0:N}, \Lambda^{(\ell-1)}) = \alpha_n^N(i) \psi_n(i) \Lambda_{ji}^{(\ell-1)} \pi_{n-1}(j)$. It also happens that equation (8.7) enforces non-negativity of Λ_{ji} , which is required for well-posedness of the algorithm. Equation (8.7) is equivalent to the estimates that were conjectured in (8.3) and (8.4).

8.3 The Viterbi Algorithm

Sometimes it may be more important to estimate the entire path of X . The Viterbi algorithm applies the properties of the HMM along with dynamic programming to find an optimal sequence $\hat{V}_{0:N} \in \mathcal{S}^{N+1}$ that maximizes the joint-posterior probability

$$\hat{V}_{0:N} = (\hat{V}_0, \dots, \hat{V}_N) \doteq \arg \max_{\vec{x} \in \mathcal{S}^{N+1}} \mathbb{P}(X_{0:N} = \vec{x}, Y_{0:N}).$$

Given the data $Y_{0:N}$, smoothing can be used to ‘look-back’ and make estimates of X_n for some $n < N$, but neither equations (8.1) or (8.2) is a joint posterior, meaning that they will not be able to tell us the posterior probability of a path $\vec{x} \in \mathcal{S}^{N+1}$. The size of our problem

would grow exponentially with N if we needed to compute the posterior distribution of X 's paths, but the Viterbi algorithm allows us to obtain the MAP estimator of X 's path with without actually calculating the posterior probabilities of all paths.

The memoryless channel of the HMM allows us to write the maximization over paths as a nested maximization,

$$\begin{aligned} \max_{\vec{x} \in \mathcal{S}^{N+1}} \mathbb{P}(X_{0:N} = \vec{x}, Y_{0:N}) &= \max_{v \in \mathcal{S}} \psi_N(v) \max_{\vec{x} \in \mathcal{S}^N} \Lambda(v|\vec{x}_{N-1}) \mathbb{P}(X_{0:N-1} = \vec{x}, Y_{0:N-1}) \\ &= \psi_N(\hat{V}_N) \max_{\vec{x} \in \mathcal{S}^N} \Lambda(\hat{V}_N|\vec{x}_{N-1}) \mathbb{P}(X_{0:N-1} = \vec{x}, Y_{0:N-1}), \quad (\dagger) \end{aligned}$$

where ψ is the likelihood and c_N is the normalizing constant, both from the forward Baum-Welch equation in (8.1). To take advantage of this nested structure, it helps to define the following recursive function,

$$\begin{aligned} \phi_0(v) &\doteq \psi_0(v) \mathbb{P}(X_0 = v) \\ \phi_n(v) &\doteq \psi_n(v) \max_x \Lambda(v|x) \phi_{n-1}(x), \quad \text{for } n = 1, 2, 3, \dots, N. \end{aligned}$$

We then place ϕ is the nested structure of (\dagger) and work backwards to obtain the optimal path,

$$\begin{aligned} \hat{V}_N &= \arg \max_v \phi_N(v) \\ \hat{V}_n &= \arg \max_v \Lambda(\hat{V}_{n+1}|v) \phi_n(v), \quad \text{for } n = N-1, N-2, \dots, 2, 1, 0 \end{aligned}$$

thus obtaining the optimal path in $O(N)$ -many computations. It would have taken $O(|\mathcal{S}|^N)$ -many computations to obtain the posterior distribution of the paths.

We are interested in the Viterbi algorithm mainly because the path of estimates returned by the filtering and smoothing may

Remark 9. *The unnormalized probabilities in ϕ quickly fall below machine precision levels, so it is better to consider a logarithmic version of Viterbi,*

$$\begin{aligned} \log \phi_0(v) &= \log \psi_0(v) + \log \mathbb{P}(X_0 = v) \\ \log \phi_n(v) &= \log \psi_n(v) + \max_x \{\log \Lambda(v|x) + \log \phi_{n-1}(x)\} \end{aligned}$$

and the use the $\log \phi_n$'s in the dynamic programming step,

$$\begin{aligned} \hat{V}_N &= \arg \max_v \log \phi_N(v) \\ \hat{V}_n &= \arg \max_v \left\{ \log \Lambda(\hat{V}_{n+1}|v) + \log \phi_n(v) \right\}, \quad \text{for } n = N-1, N-2, \dots, 2, 1, 0. \end{aligned}$$

Chapter 9

The Particle Filter

Monte Carlo methods have become the most common way to compute quantities from HMMs –and with good reason; they are in fact a fast and effective way to obtain consistent estimates. In particular, the particle filter is used to approximate filtering expectations. There are similar methods that exploit Bayes formula in obtaining samples from an HMM, but ‘particle filtering’ implies that sequential Monte Carlo (SIS) and Sampling-Importance-Resampling (SIR) are applied to the specified HMM.

9.1 The Particle Filter

Suppose that X_n is an unobserved Markov chain taking values in a state-space denoted by \mathcal{S} . Let Λ denote X_n ’s kernel of transition densities so that

$$\frac{d}{dx}\mathbb{P}(X_{n+1} \leq x) = \int \Lambda(x|v)\mathbb{P}(X_n \in dv)$$

for any $x \in \mathcal{S}$, and $\frac{d}{dx}\mathbb{P}(X_0 \leq x) = p_0(x)$. Let the observed process Y_n be a nonlinear function of X_n ,

$$Y_n = h(X_n) + W_n$$

where W_n is an iid Gaussian random variable with mean zero and variance $\gamma^2 > 0$. In this case, the filter is easily shown to be a density function, given recursively as,

$$\pi_{n+1}(x) = \frac{1}{c_{n+1}}\psi_{n+1}(x) \int \Lambda(x|v)\pi_n(v)$$

where c_{n+1} is a normalizing constant, and ψ_{n+1} is a likelihood function

$$\psi_{n+1}(x) = \exp \left\{ -\frac{1}{2} \left(\frac{Y_{n+1} - h(x)}{\gamma} \right)^2 \right\},$$

but some kind quadrature grid would need to be established over \mathcal{S} if we were to use this recursive expression. An alternative is to use particles.

9.1.1 Sequential Importance Sampling (SIS)

Ideally, we would be able to sample directly from the filtering distribution to obtain a Monte Carlo estimate,

$$\frac{1}{P} \sum_{\ell=1}^P g(x_n^\ell) \approx \mathbb{E}[g(X_n) | \mathcal{F}_n^Y], \quad \text{for } P \text{ large,}$$

where $x_n^\ell \sim iid \pi_n(x)$. However, difficulties in computing π_n also make it difficult to obtain samples. However, with relative ease we can sequentially obtain samples from X' 's unconditional distribution and then assign them weights in such a way that approximates the filter.

For $\ell = 1, 2, 3, 4, \dots$, each particle is a path $x_{0:N}^\ell$ that is generated according to the unconditional distribution,

$$\begin{aligned} x_0^\ell &\sim p_0(\cdot) \\ x_n^\ell &\sim \Lambda(\cdot | x_{n-1}^\ell) \quad \text{for } n = 1, 2, 3, \dots, N. \end{aligned}$$

Then for P -many particles and any integrable function g , the strong law of large numbers tells us that

$$\frac{1}{P} \sum_{\ell=1}^P g(x_{0:N}^\ell) \rightarrow \mathbb{E}[g(X_{0:N})]$$

almost surely as $P \rightarrow \infty$.

Given $Y_{0:n}$, let ω_n^ℓ denote the *importance weight* of a particle. We define ω_n^ℓ to be proportional to the likelihood of the ℓ th particle's path, which we can write recursively as the product of its old weight and a likelihood function:

$$\omega_n^\ell = \frac{1}{c_n} \mathbb{P}(Y_{0:n} | X_{0:n} = x_{0:n}^\ell) = \frac{1}{c_n} \psi_n(x_n^\ell) \omega_{n-1}^\ell \quad \text{for } n = 0, 1, 2, 3, \dots, N$$

with the convention that $\omega_{-1}^\ell \equiv 1$, and c_n is a normalizing constant

$$c_n = \sum_{\ell=1}^P \psi_n(x_n^\ell) \omega_{n-1}^\ell.$$

Then the filtering expectation of an integrable function $g(X_N)$ can be consistently approximated with the weighted particles,

$$\sum_{\ell=1}^P g(x_N^\ell) \omega_N^\ell = \frac{\frac{1}{P} \sum_{\ell=1}^P g(x_N^\ell) \mathbb{P}(Y_{0:N} | X_{0:N} = x_{0:N}^\ell)}{\frac{1}{P} \sum_{\ell=1}^P \mathbb{P}(Y_{0:N} | X_{0:N} = x_{0:N}^\ell)}$$

$$\rightarrow \frac{\mathbb{E} \left[g(\tilde{X}_N) \mathbb{P}(Y_{0:N} | \tilde{X}_{0:N}) \middle| \mathcal{F}_N^Y \right]}{\mathbb{E} \left[\mathbb{P}(Y_{0:N} | \tilde{X}_{0:N}) \middle| \mathcal{F}_N^Y \right]} = \mathbb{E}[g(X_N) | \mathcal{F}_N^Y]$$

almost surely as $P \rightarrow \infty$ by SLLN, where $\tilde{X}_{0:N}$ is a random variable with distribution (p_0, Λ) and independent from $(X_{0:N}, Y_{0:N})$.

9.1.2 Sampling Importance Resampling (SIR)

Our estimation of $\mathbb{E}[g(X_N) | \mathcal{F}_N^Y]$ becomes a particle filter when SIR is used along with SIS. SIR essentially invokes a bootstrap on the samples $\{x_{0:n}^\ell\}_\ell$ at time n . This procedure will reallocate our sampling resources onto particles that are more likely to be close to the true signal. When invoked, SIR does the following:

Algorithm 1. SIR Bootstrap Procedure.

for $\ell = 1, \dots, P$ **do**

sample a random variable $x_n^{\ell, sir}$ from $\{x_n^1, \dots, x_n^P\}$ according to $\{\omega_n^1, \dots, \omega_n^P\}$.

end for

$\{x_n^1, \dots, x_n^P\} \leftarrow \{x_n^{1, sir}, \dots, x_n^{P, sir}\}$.

$\{\omega_n^1, \dots, \omega_n^P\} \leftarrow \{1/P, \dots, 1/P\}$.

The common criterion for invoking SIR can be related to an entropy approximation of the particle distribution. At any time n prior to when SIR has been performed, the entropy is defined as

$$\mathcal{E}_n = - \sum_{\ell} \omega_n^\ell \log \omega_n^\ell \geq - \log \left(\sum_{\ell} (\omega_n^\ell)^2 \right) > 0$$

and so maximizing the entropy of the particle distribution is approximately the same as minimizing the sum of squared posterior weights. Therefore, the criterion is to invoke SIR whenever the number of important particles is less than some threshold $\delta \in [1, P]$:

$$\text{if } \frac{1}{\sum_{\ell} (\omega_n^\ell)^2} \leq \delta, \quad \text{then invoke SIR.}$$

Even after SIR has been incorporated, our approximation is still consistent with the nonlinear filter:

Theorem 9.1.1. *For any bounded function $g(x)$,*

$$\frac{1}{P} \sum_{\ell=1}^P g(x_n^{\ell, sir}) \rightarrow \mathbb{E}[g(X_n) | \mathcal{F}_n^Y]$$

in L^2 as $P \rightarrow \infty$, and $x_n^{\ell, sir}$ and $x_n^{\ell', sir}$ are asymptotically independent for any $\ell \neq \ell'$.

Proof. (taken from section 9.2 of [11]) Let $\{x_n^\ell\}_{\ell \leq P}$ be the set SIS samples that were in use prior to SIR. The post-SIR estimator can be written as a sum of the old samples:

$$\frac{1}{P} \sum_{\ell} g(x_n^{\ell, sir}) = \frac{1}{P} \sum_{\ell} g(x_n^\ell) \cdot \tau_\ell$$

where τ_ℓ is the number of times the x_n^ℓ was resampled, $\tau_\ell = \sum_{r=1}^P \mathbf{1}_{\{x_n^{r, sir} = x_n^\ell\}}$. Taking conditional expectations, we have $E[\tau_\ell | \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P}] = P \cdot \omega_n^\ell$ and the conditional expectation of the estimator is

$$\begin{aligned} E \left[\frac{1}{P} \sum_{\ell} g(x_n^{\ell, sir}) \middle| \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P} \right] &= \frac{1}{P} \sum_{\ell} g(x_n^\ell) E[\tau_\ell | \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P}] \\ &= \sum_{\ell} g(x_n^\ell) \omega_n^\ell \xrightarrow{a.s.} \mathbb{E} \left[g(X_n) \middle| \mathcal{F}_n^Y \right], \quad \text{as } P \rightarrow \infty. \end{aligned}$$

From here we take expectations of both sides and use dominated convergence to equate the limit to show L^2 convergence,

$$E \left| \frac{1}{P} \sum_{\ell=1}^P g(x_n^{\ell, sir}) - \mathbb{E} \left[g(X_n) \middle| \mathcal{F}_n^Y \right] \right|^2 \rightarrow 0$$

as $P \rightarrow \infty$.

Now consider another bounded function $f(x)$,

$$\begin{aligned} \mathbb{E} \left[g(x_n^{\ell, sir}) f(x_n^{\ell', sir}) \middle| \mathcal{F}_n^Y \right] &= \mathbb{E} \left[\mathbb{E} \left[g(x_n^{\ell, sir}) f(x_n^{\ell', sir}) \middle| \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P} \right] \middle| \mathcal{F}_n^Y \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[g(x_n^{\ell, sir}) \middle| \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P} \right] \mathbb{E} \left[f(x_n^{\ell', sir}) \middle| \mathcal{F}_n^Y \vee \{x_n^r\}_{r \leq P} \right] \middle| \mathcal{F}_n^Y \right] \\ &= \mathbb{E} \left[\left(\sum_{\ell} g(x_n^\ell) \omega_n^\ell \right) \left(\sum_{\ell'} f(x_n^{\ell'}) \omega_n^{\ell'} \right) \middle| \mathcal{F}_n^Y \right] \rightarrow \mathbb{E}[g(X_n) | \mathcal{F}_n^Y] \cdot \mathbb{E}[f(X_n) | \mathcal{F}_n^Y], \end{aligned}$$

as $P \rightarrow \infty$. So we've found that

$$\mathbb{E} \left[g(x_n^{\ell, sir}) f(x_n^{\ell', sir}) \middle| \mathcal{F}_n^Y \right] \sim \mathbb{E} \left[g(x_n^{\ell, sir}) \middle| \mathcal{F}_n^Y \right] \mathbb{E} \left[f(x_n^{\ell', sir}) \middle| \mathcal{F}_n^Y \right]$$

for $\ell \neq \ell'$ and P large. Therefore, $x_n^{\ell, sir}$ and $x_n^{\ell', sir}$ are asymptotically independent. \square

Variance Reduction

For any bounded function $g(x)$, the principle of conditional Monte Carlo tells us that SIR estimator will have greater variance than the SIS estimator,

$$\begin{aligned}
& \text{var} \left(\frac{1}{P} \sum_{\ell'} g(x_n^{\ell', \text{sir}}) \right) \\
&= \text{var} \left(\frac{1}{P} \sum_{\ell'} g(x_n^{\ell', \text{sir}}) \middle| \{x_n^\ell, \omega_n^\ell\}_\ell \right) + \text{var} \left(\mathbb{E} \left[\frac{1}{P} \sum_{\ell'} g(x_n^{\ell', \text{sir}}) \middle| \{x_n^\ell, \omega_n^\ell\}_\ell \right] \right) \\
&\geq \text{var} \left(\mathbb{E} \left[\frac{1}{P} \sum_{\ell'} g(x_n^{\ell', \text{sir}}) \middle| \{x_n^\ell, \omega_n^\ell\}_\ell \right] \right) = \text{var} \left(\sum_\ell g(x_n^\ell) \omega_n^\ell \right)
\end{aligned}$$

and so the estimator $\frac{1}{P} \sum_{\ell'} g(x_n^{\ell', \text{sir}})$ may not be preferable to $\sum_\ell g(x_n^\ell) \omega_n^\ell$. However, it follows from the proof of theorem 9.1.1 that

$$\text{var}(g(X_n) | \mathcal{F}_n^Y) \sim \text{var}(g(x_n^{\ell, \text{sir}}) | \mathcal{F}_n^Y), \quad \text{for } P \text{ large,}$$

and we see a reduction in the overall variance of the particles if we write down the law of total variance,

$$\begin{aligned}
\text{var}(g(x_n^\ell)) &= \text{var}(g(X_n)) = \text{var} \left(g(X_n) \middle| \mathcal{F}_n^Y \right) + \underbrace{\text{var} \left(\mathbb{E}[g(X_n) | \mathcal{F}_n^Y] \right)}_{>0} \\
&> \text{var} \left(g(X_n) \middle| \mathcal{F}_n^Y \right) \sim \text{var}(g(x_n^{\ell, \text{sir}})).
\end{aligned}$$

This reduction can be quite significant if $g(X_n)$ has a broad range. The rates are difficult to show, but by invoking SIR we obtain estimates of $\mathbb{E}[g(X_n) | \mathcal{F}_n^Y]$ that will converge faster as $P \rightarrow \infty$. This brief subsection has not attempted any proof; we have not computed any comparison of convergence rates.

9.2 Examples

In this section we present some examples to demonstrate the particle filter's uses.

9.2.1 Particle Fiter for Heston Model

Consider a Heston model with time-dependent coefficients

$$\begin{aligned} dY_t &= \left(\mu - \frac{1}{2} X_t \right) dt + \sqrt{X_t} \left(\rho dB_t + \sqrt{1 - \rho^2} dW_t \right) \\ dX_t &= \nu(\bar{X} - X_t)dt + \gamma\sqrt{X_t}dB_t \end{aligned}$$

where Y_t is the log-price of an equity, $\sqrt{X_t}$ is the volatility, $\rho \in [-1, 1]$ is the correlation parameter, and (W_t, B_t) are a pair of independent Wiener processes. The observed log-prices on equities and indices is not available in continuum. Instead, there is a discrete set sequence $(t_n)_{n=0,1,2,\dots}$ consisting of times at which quotes on the equity or index are given,

$$Y_n \doteq Y_{t_n}, \quad \text{for } n = 0, 1, 2, 3, 4, \dots$$

We denote the time step between the n th and $(n+1)$ th observations $\Delta t_n = t_{n+1} - t_n$. By considering the Stratonovich/Itô integral transform

$$\int_{t_n}^{t_{n+1}} \gamma_s \sqrt{X_s} \circ dB_s = \frac{1}{2} \int_{t_n}^{t_{n+1}} \gamma_s^2 ds + \int_{t_n}^{t_{n+1}} \gamma_s \sqrt{X_s} dB_s$$

and letting $X_n \doteq X_{t_n}$, we will find it useful to work with the following implicit discretization of the Stratonovich form of the Heston model,

$$\begin{aligned} Y_{n+1} &= Y_n + \left(\mu - \frac{1}{2} X_n \right) \Delta t_n + \sqrt{X_n} \left(\rho \Delta B_n + \sqrt{1 - \rho^2} \Delta W_n \right) \\ X_{n+1} &= X_n (1 - \nu \Delta t_n) + \left(\nu \bar{X} - \frac{\gamma^2}{2} \right) \Delta t_n + \gamma \sqrt{X_{n+1}} \Delta B_n \quad (*) \end{aligned}$$

where ΔB_n and ΔW_n are increments of independent Wiener processes (i.e. $\Delta B_n \doteq B_{t_{n+1}} - B_{t_n} \sim N(0, \Delta t_n)$ and $\Delta W_n \doteq W_{t_{n+1}} - W_{t_n} \sim N(0, \Delta t_n)$). We take $\sqrt{X_{n+1}}$ to be the root of equation (*) which can be obtained through the quadratic equation (see Alfonsi [2]),

$$\sqrt{X_{n+1}} = \frac{1}{2} \left\{ \gamma \Delta B_n \pm \sqrt{\gamma^2 \Delta B_n^2 + 4D} \right\}$$

$$\text{where} \quad D = (1 - \nu \Delta t_n) X_n + \left(\nu \bar{X} - \frac{\gamma^2}{2} \right) \Delta t_n.$$

Provided that $\frac{1}{n} = \Delta t \leq \frac{1}{\nu}$ and $\gamma^2 \leq 2\nu\bar{X}$, this implicit scheme is effective because it is mean-reverting and preserves positivity in X^n . With this scheme we can generate particles $\{x_n^\ell\}_{n,\ell}$ and approximate the nonlinear filter.

9.2.2 Rao-Blackwellization

Let θ_n be a hidden Markov chain with transition probabilities Λ , and let X_n be another hidden Markov process given by the following recursion,

$$X_n = a(\theta_n)X_{n-1} + \sigma(\theta_n)B_n$$

with $B_n \sim iidN(0, 1)$ (to be clear, $B_n \perp \theta_n$), and Gaussian initial distribution $p_0(x)$. Let the observations process be defined discretely as

$$Y_n = h(\theta_n)X_n + \gamma(\theta_n)W_n$$

where W_k are $iidN(0, 1)$ (to be clear $B_n \perp W_n$ and θ_n independent of W_n), and $\gamma(\cdot) > 0$. In this case we can use particles to marginalize θ_n ,

$$\begin{aligned} \theta_0^\ell &\sim p_0 \\ \theta_n^\ell &\sim \Lambda(\cdot | \theta_{n-1}^\ell) \quad \text{for } n > 0, \end{aligned}$$

and then for each particle we can compute the marginal Kalman filter,

$$\begin{aligned} G_n^\ell &= \frac{h(\theta_n^\ell)\Sigma_{n-1}^\ell}{h^2(\theta_n^\ell)(a^2(\theta_n^\ell)\Sigma_{n-1}^\ell + \sigma^2(\theta_n^\ell)) + \gamma^2(\theta_n^\ell)} \\ \hat{X}_n^\ell &= a(\theta_n^\ell)\hat{X}_{n-1}^\ell + G_n^\ell(Y_n - h(\theta_n^\ell)a(\theta_n^\ell)\hat{X}_{n-1}^\ell) \\ \Sigma_n^\ell &= (1 - G_n^\ell h(\theta_n^\ell)) (a^2(\theta_n^\ell)\Sigma_{n-1}^\ell + \sigma^2(\theta_n^\ell)) \end{aligned}$$

where we have defined $\hat{X}_n^\ell \doteq \mathbb{E}[X_n | \mathcal{F}_n^Y \vee \{\theta_{0:n}^\ell\}]$ and $\Sigma_n^\ell \doteq \mathbb{E}[(X_n - \hat{X}_n^\ell)^2 | \theta_{0:n}^\ell]$ (see [28, 25] for more on the Kalman Filter). Conditioned on $\mathcal{F}_{n-1}^Y \vee \{\theta_{0:n}^\ell\}$, Y_n is normal with mean and variance

$$\begin{aligned} \mu_{n|n-1}^\ell &\doteq \mathbb{E}[Y_n | \mathcal{F}_{n-1}^Y \vee \{\theta_{0:n}^\ell\}] = h(\theta_n^\ell)a(\theta_n^\ell)\hat{X}_{n-1}^\ell \\ v_{n|n-1}^\ell &\doteq \text{var}\left(Y_n \middle| \mathcal{F}_{n-1}^Y \vee \{\theta_{0:n}^\ell\}\right) = h^2(\theta_n^\ell) \left(a^2(\theta_n^\ell)\Sigma_{n-1}^\ell + \sigma^2(\theta_n^\ell) \right) + \gamma^2(\theta_n^\ell) \end{aligned}$$

with the convention $\mu_{0:-1}^\ell = h(\theta_0^\ell)a(\theta_0^\ell)\mathbb{E}X_0$ and $v_{0:-1}^\ell = h^2(\theta_0^\ell)a^2(\theta_0^\ell)\text{var}(X_0) + \gamma^2(\theta_0^\ell)$. For any particle $\theta_{0:n}^\ell$, the unnormalized importance weights are updated as follows:

$$\tilde{\omega}_n^\ell \doteq \mathbb{P}(Y_{0:n} | \theta_{0:n} = \theta_{0:n}^\ell) = \frac{\exp\left\{-\frac{1}{2}\left(\frac{Y_n - \mu_{n|n-1}^\ell}{\sqrt{v_{n|n-1}^\ell}}\right)^2\right\}}{\sqrt{v_{n|n-1}^\ell}} \times \tilde{\omega}_{n-1}^\ell$$

with the convention that $\tilde{\omega}_{-1}^\ell \equiv 1$. For P -many particles, the unnormalized Rao-Blackwellized filter is an approximation to the unnormalized filter,

$$\phi_n^*[\theta] \doteq \frac{1}{P} \sum_{\ell} \theta_n^\ell \tilde{\omega}_n^\ell = \frac{1}{P} \sum_{\ell} \theta_n^\ell \mathbb{P}(Y_{0:n} | \theta_{0:n} = \theta_{0:n}^\ell) \approx \mathbb{E}[\tilde{\theta}_n \mathbb{P}(Y_{0:n} | \tilde{\theta}_{0:n}) | \mathcal{F}_n^Y] \doteq \tilde{\mathbb{E}}[\theta_n | \mathcal{F}_n^Y]$$

where $\tilde{\theta}$ is a copy of θ that is independent of (Y, X, θ) .

The Rao-Blackwell theorem says the following:

Theorem 9.2.1. Rao-Blackwell. *Given $Z_{0:n}$, let $\hat{\beta}$ be an estimator of a parameter β , and T a sufficient statistic for β . Then the estimator $\hat{\beta}^* = \mathbb{E}[\hat{\beta} | T(Z_{0:n})]$ is at least as good in terms of MSE*

$$\mathbb{E}(\hat{\beta}^* - \beta)^2 \leq \mathbb{E}(\hat{\beta} - \beta)^2$$

for all β in the parameter space.

Now, suppose that for each ℓ we generate particles $\{x_{0:n}^{\ell', \ell}\}_{\ell' \leq P'}$ instead of computing Kalman filters, and then using these particles we compute another estimator of the unnormalized filtering expectation,

$$\phi_n[\theta] \doteq \frac{1}{P} \sum_{\ell} \theta_n^\ell \underbrace{\frac{1}{P'} \sum_{\ell'} \mathbb{P}(Y_{0:n} | X_{0:n} = x_{0:n}^{\ell', \ell}, \theta_{0:n} = \theta_{0:n}^\ell)}_{\approx \mathbb{P}(Y_{0:n} | \theta_{0:n} = \theta_{0:n}^\ell)}.$$

But given $\theta_{0:n}^\ell$, the marginal Kalman filter has allowed us to compute the likelihood without approximation. Therefore, a sufficient statistic for $\tilde{\mathbb{E}}[\theta_n | \mathcal{F}_n^Y]$ is $\mathcal{T}_n \doteq (\theta_n^\ell, \tilde{\omega}_n^\ell)_{\ell \leq P}$, and we have

$$\begin{aligned} \tilde{\omega}_n^\ell &= \mathbb{P}(Y_{0:n} | \theta_{0:n} = \theta_{0:n}^\ell) = \int \mathbb{P}(Y_{0:n} | X_{0:n} = x_{0:n}, \theta_{0:n} = \theta_{0:n}^\ell) \mathbb{P}(X_{0:n} \in dx_{0:n} | \theta_{0:n} = \theta_{0:n}^\ell) \\ &= \mathbb{E} \left[\mathbb{P}(Y_{0:n} | X_{0:n} = x_{0:n}^{\ell', \ell}, \theta_{0:n} = \theta_{0:n}^\ell) \middle| \mathcal{F}_n^Y \vee \{\theta_{0:n}^\ell\} \right] \end{aligned}$$

for all $\ell' \leq P'$, and from here it is easy to see that the expectation of $\phi_n[\theta]$ given $\mathcal{F}_n^Y \vee \mathcal{T}_n$ is the Rao-Blackwellized estimator,

$$\mathbb{E}[\phi_n[\theta] | \mathcal{F}_n^Y \vee \mathcal{T}_n] = \mathbb{E} \left[\mathbb{E}[\phi_n[\theta] | \mathcal{F}_n^Y \vee \{\theta_{0:n}^\ell\}_{\ell}] \middle| \mathcal{F}_n^Y \vee \mathcal{T}_n \right] = \frac{1}{P} \sum_{\ell} \theta_n^\ell \tilde{\omega}_n^\ell = \phi_n^*[\theta].$$

Therefore, by the Rao-Blackwell theorem we know that $\phi_n^*[\theta]$ has less or equal MSE to a particle filter computed without marginal Kalman filters,

$$\mathbb{E} \left[\left(\phi_n^*[\theta] - \tilde{\mathbb{E}}[\theta_n | \mathcal{F}_n^Y] \right)^2 \middle| \mathcal{F}_n^Y \right] \leq \mathbb{E} \left[\left(\phi_n[\theta] - \tilde{\mathbb{E}}[\theta_n | \mathcal{F}_n^Y] \right)^2 \middle| \mathcal{F}_n^Y \right],$$

with the advantage that the Rao-Blackwellized filter requires particles to be simulated across a domain of fewer dimensions.

Chapter 10

Stability, Lyapunov Exponents, and Ergodic Theory for Finite-State Filters

A filter is said to be ‘stable’ if it has the ability to asymptotically recover from an erroneous initial distribution. In other words, assuming that all parts of the HMM are estimated correctly, but with the exception of the initial distribution which is incorrect, a stable filter will ‘forget’ the false assumptions as the initial data falls farther and farther into the past. It can be advantageous to work with a stable filter for a number of reasons, one being that parameter estimation algorithms for models with stable filters do not need to put as much emphasis on estimating the initial condition. Rates at which stability takes effect can also be estimated, as these rates are shown to be given by Lyapunov exponents or a spectral gap in the filter’s generator, and while the exponents are generally not explicitly computable, there are ways to make estimates. Finally, it is possible that a stable filter for an ergodic state variable may also have an ergodic theorem, a property that may be also be useful for parameter estimation.

10.1 Main Ideas and Their History

Consider an HMM (X_t, Y_t) where X_t is a hidden Markov process taking values in a state-space \mathcal{S} with initial distribution $\nu : \mathcal{S} \rightarrow [0, 1]$, and where observations are made on the process Y_t which we assume to be given by a function of X_t plus a noise. A filtering measure can be computed using the initial condition ν , and we denote it as

$$\pi_t^\nu(A) = \mathbb{P}^\nu(X_t \in A | \sigma\{Y_s : s \leq t\}) \quad \text{for all Borel sets } A \subset \mathcal{S}.$$

Definition 10.1.1. *Let $\tilde{\nu}$ be another probability measure on \mathcal{S} , and let $\pi_t^{\tilde{\nu}}$ denote the filter computed with $\tilde{\nu}$ as the initial distribution of X . The filter is said to be asymptotically*

stable if

$$\lim_{t \rightarrow \infty} \mathbb{E} |\pi_t^\nu - \pi_t^{\tilde{\nu}}|_{TV} = 0$$

where $|\cdot|_{TV}$ denotes the total-variation norm.¹

Intuitively, the filter should be stable if the model fits one of the following descriptions:

- the signal is ergodic
- or the observations are sufficiently informative, making old information obsolete (e.g. very low noise and $h(\cdot)$ is one-to-one),

but it is difficult to prove stability results even for these basic cases, and a general theory has yet to be developed.

Stability results for finite-state Markov chain signals are well-known along their rates, and their ergodic theory, and it has been known since the 1960's that Kalman and Kalman-Bucy filters are stable when the signal is ergodic. The fundamental way of showing filter stability for these filters is to identify the equilibrium of the posterior covariance from its Ricatti equations, use it to show that the gain matrix also approaches an equilibrium, and then verify that the dependence on the initial condition fades with time. Example 10.1.1 shows how this is done for general Kalman-Bucy filters with constant coefficients, thus showing that observations in linear Gaussian models are indeed sufficiently informative since the signal is not assumed to be ergodic.

Example 10.1.1. (Stability of Kalman-Bucy Filter with Constant Coefficients). Consider the following linear system,

$$\begin{aligned} dX_t &= aX_t dt + \sigma dB_t \\ dY_t &= hX_t dt + \gamma dW_t \end{aligned}$$

where $B_t \perp W_t$. Applying the Kalman-Bucy filter, we have

$$\begin{aligned} d\hat{X}_t &= \left(a - \frac{h^2}{\gamma^2} \Sigma_t \right) \hat{X}_t dt + \frac{h}{\gamma^2} \Sigma_t dY_t \\ \frac{d}{dt} \Sigma_t &= 2a\Sigma_t - \frac{h^2}{\gamma^2} \Sigma_t^2 + \sigma^2. \end{aligned}$$

The solution to the Ricatti equation can be written explicitly as follows,

$$\Sigma_t = \frac{\alpha_1 - K\alpha_2 \exp \left\{ \frac{h^2}{\gamma^2} (\alpha_2 - \alpha_1) t \right\}}{1 - K \exp \left\{ \frac{h^2}{\gamma^2} (\alpha_2 - \alpha_1) t \right\}}$$

¹The total variation norm of the difference between two probability measure p and q is $|p - q|_{TV} = \sup\{|p(A) - q(A)| : A \in \mathcal{B}(\mathbb{R})\}$ where $\mathcal{B}(\mathbb{R})$ denotes the space of Borel-measurable subset of \mathbb{R} .

where

$$\alpha_1 = h^{-2} \left(a\gamma^2 - \gamma\sqrt{a^2\gamma^2 + h^2\sigma^2} \right), \quad \alpha_2 = h^{-2} \left(a\gamma^2 + \gamma\sqrt{a^2\gamma^2 + h^2\sigma^2} \right)$$

with $K = \frac{\Sigma_0 - \alpha_1}{\Sigma_0 - \alpha_2}$. Asymptotically, we have $\Sigma_t \sim \alpha_2$ and the filtering expectation is approximately,

$$\hat{X}_t \sim \hat{X}_0 e^{-\beta t} + \frac{h\alpha_2}{\gamma^2} \int_0^t e^{-\beta(t-s)} dY_s$$

where $\beta = \frac{1}{\gamma} \sqrt{a^2\gamma^2 + h^2\sigma^2}$. This shows that the Kalman-Bucy filter will forget any initial condition on X_0 as $t \rightarrow \infty$, thus showing that the filter is stable.

General results for nonlinear filters with ergodic states were identified by Kunita in 1971 [31], but a key step in his proof is wrong. His proof essentially said the following:

Assume that $\nu \ll \tilde{\nu}$. Then for any test function $g(x)$,

$$\begin{aligned} \int_S g(x) d\pi_t^\nu(x) &= \mathbb{E}[g(X_t) | \mathcal{F}_t^Y] = \frac{\mathbb{E}^{\tilde{\nu}} \left[g(X_t) \frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]}{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]} \\ &= \mathbb{E}^{\tilde{\nu}} \left[g(X_t) \frac{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \vee \{X_t\} \right]}{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]} \middle| \mathcal{F}_t^Y \right] = \int_S g(x) \frac{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \vee \{X_t\} \right]}{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]} d\pi_t^{\tilde{\nu}}(x) \end{aligned}$$

where the denominator is strictly positive a.s. because $\nu \ll \tilde{\nu}$. From this we see that the $\pi^\nu \ll \pi^{\tilde{\nu}}$ with Radon-Nykodym derivative

$$\frac{d\pi_t^\nu}{d\pi_t^{\tilde{\nu}}}(x) = \frac{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \vee \{X_t = x\} \right]}{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]}$$

with \mathbb{P} -a.s. From the existence of the Radon-Nykodym derivative, the TV-norm is equivalent \mathbb{P} -a.s. to the following

$$\begin{aligned} \|\pi_t^\nu - \pi_t^{\tilde{\nu}}\|_{TV} &= \int_S \left| \frac{d\pi_t^\nu}{d\pi_t^{\tilde{\nu}}}(x) - 1 \right| d\pi_t^{\tilde{\nu}}(x) \\ &= \frac{\mathbb{E}^{\tilde{\nu}} \left[\left| \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \vee \{X_t\} \right] - \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right] \right| \middle| \mathcal{F}_t^Y \right]}{\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right]} \end{aligned} \quad (10.1)$$

but because of the Markov property, we realize that the distribution of X_0 given $\mathcal{F}_t^Y \vee \{X_t\}$ is the same regardless of whether information is added regarding the future. Therefore, we have

$$\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \vee \{X_t\} \right] = \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t,\infty)}^X \right]$$

where $\mathcal{F}_{[t,\infty)}^X$ denotes the tail- σ -field generated by $\{X_s : s \geq t\}$. Combining the tail- σ -field measure with the numerator in (10.1) we have

$$\begin{aligned} \mathbb{E} \|\pi_t^\nu - \pi_t^{\tilde{\nu}}\|_{TV} &= \mathbb{E}^{\tilde{\nu}} \left[\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right] \|\pi_t^\nu - \pi_t^{\tilde{\nu}}\|_{TV} \right] \\ &= \mathbb{E}^{\tilde{\nu}} \left| \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t,\infty)}^X \right] - \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right] \right| \end{aligned}$$

and taking limits we have

$$\lim_{t \rightarrow \infty} \mathbb{E} \|\pi_t^\nu - \pi_t^{\tilde{\nu}}\|_{TV} = \mathbb{E}^{\tilde{\nu}} \left| \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \bigcap_{t \geq 0} \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t,\infty)}^X \right] - \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_t^Y \right] \right|$$

which suggests that the filters are stable if and only if

$$\mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \bigcap_{t \geq 0} \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t,\infty)}^X \right] = \mathbb{E}^{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(X_0) \middle| \mathcal{F}_\infty^Y \right]. \quad (10.2)$$

To this point, every step is correct, but the error made by Kunita was in assuming that the limits of these filtrations were equal,

$$\bigcap_{t \geq 0} \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t,\infty)}^X \stackrel{?}{=} \mathcal{F}_\infty^Y, \quad (10.3)$$

but there has since come a counter-example to equation (10.3).

10.1.1 The Counter Example

Baxendale, Chigansky and Lipster [8] presented an example to demonstrate when (10.3) fails. Let X_t be a Markov chain taking values in $\mathcal{S} = \{1, 2, 3, 4\}$ with transition intensities

$$\Lambda = \begin{bmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{bmatrix}.$$

Clearly, all state communicate, and X is an ergodic Markov process with invariant measure $\mu = (1, 1, 1, 1)/4$. Let $h(x) = \mathbf{1}_{x=1} + \mathbf{1}_{x=3}$, and consider the observations model

$$Y_t = h(X_t),$$

which is a degenerate noise model. The following lemma was proven in [8],

Lemma 10.1.1. *For this example, the limit of the filtrations in equation (10.3) is false,*

$$\bigcap_{t \geq 0} \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t, \infty)}^X \supsetneq \mathcal{F}_\infty^Y.$$

Proof. It suffices to show that X_0 is $\mathcal{F}_\infty^Y \vee \mathcal{F}_{[t, \infty)}^X$ -measurable, but not measurable with respect to \mathcal{F}_∞^Y . The Markov chain X only admits cycles in the following order,

$$\dots\dots\{3\} \rightarrow \{4\} \rightarrow \{1\} \rightarrow \{2\} \rightarrow \dots,$$

and therefore we can recover X_0 given \mathcal{F}_t^Y and X_t for any $t > 0$ (i.e. because we know X_t we can look backwards and deduce the path of X by looking at what times Y has jumped). Now, because $\mathcal{F}_t^Y \vee \{X_t\} \subset \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t, \infty)}^X$, we have

$$X_0 \in \bigcap_{t \geq 0} \mathcal{F}_\infty^Y \vee \mathcal{F}_{[t, \infty)}^X.$$

Next, denote the times at which Y jumps with the sequence $\{\tau_i\}_{i \geq 1}$ (τ_i is time of Y 's i th jump). It is not hard to verify that τ_i is independent of (X_0, Y_0) and the the following filtrations are equal,

$$\mathcal{F}_t^Y = \bigvee_{i \geq 1} \{\tau_i \leq t\} \vee \{Y_0\}$$

and so for any $t > 0$ we have

$$\begin{aligned} \mathbb{P}(X_0 = 1 | \mathcal{F}_t^Y) &= \mathbb{P}\left(X_t = 1 \middle| \bigvee_{i \geq 1} \{\tau_i \leq t\} \vee \{Y_0\}\right) = \mathbb{P}(X_0 = 1 | Y_0) \\ &= \frac{\mathbb{P}(X_0 = 1)}{\mathbb{P}(X_0 = 1) + \mathbb{P}(X_0 = 3)} Y_0 \neq \mathbf{1}_{X_t=1}. \end{aligned}$$

Since this posterior holds for any $t > 0$, we must have

$$\mathbb{P}(X_0 = 1 | \mathcal{F}_\infty^Y) \neq \mathbf{1}_{X_t=1}$$

which means that $X_0 \notin \mathcal{F}_\infty^Y$. □

10.2 Stability for Markov Chain Models

In this section we present some of the results in the paper by Atar and Zeitouni [4]. In particular, we present their proof of the stability-rate for discrete-time filtering problems where the state variable is an ergodic finite-state Markov chain. Their paper also presents the analogous results for continuous time, as well as some other results regarding the low-noise case.

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and let $n = 0, 1, 2, 3 \dots$ denote time. Suppose that X_n is an unobserved Markov chain taking values in a finite state-space $\mathcal{S} = \{x_1, \dots, x_d\}$. Let the matrix Λ contain X_n 's transition probabilities so that

$$\mathbb{P}(X_{n+1} = x_i) = \sum_j \Lambda_{ji} \mathbb{P}(X_n = x_j)$$

for any $i, j \leq d$, and $\mathbb{P}(X_0 = x_i) = \nu_i$. Suppose further that X_n is recurrent with invariant law μ such that

$$(\Lambda^*)^n \nu \rightarrow \mu$$

as $n \rightarrow \infty$. We assume that X_n is ergodic, which can be the case if and only if Λ is of primitive order k (i.e. there exists $k < \infty$ such that $\Lambda_{ji}^n > 0$ for all $i, j \leq d$ and for all $n \geq k$). Let the observed process Y_n be a nonlinear function of X_n ,

$$Y_n = h(X_n) + W_n$$

where W_n is an iid Gaussian random variable with mean zero and variance $\gamma^2 > 0$. The filtering mass computed with μ as its initial condition is denoted with π_n^ν and is given recursively by

$$\pi_{n+1}^\nu = \frac{1}{c_{n+1}} \psi_{n+1} \Lambda^* \pi_n^\nu$$

where c_{n+1} is a normalizing constant (dependent on μ), and ψ_{n+1} is a diagonal matrix of likelihood functions

$$\psi_{n+1} = \begin{bmatrix} e^{-\frac{1}{2} \left(\frac{Y_{n+1} - h(x_1)}{\gamma} \right)^2} & 0 & \dots & 0 \\ 0 & e^{-\frac{1}{2} \left(\frac{Y_{n+1} - h(x_2)}{\gamma} \right)^2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{-\frac{1}{2} \left(\frac{Y_{n+1} - h(x_d)}{\gamma} \right)^2} \end{bmatrix}.$$

Stability in this case means that for any other measure $\tilde{\nu} : \mathcal{S} \rightarrow [0, 1]$, we have

$$\|\pi_n^\nu - \pi_n^{\tilde{\nu}}\| \rightarrow 0$$

as $n \rightarrow \infty$, where $\|\cdot\|$ denotes the Euclidean norm on \mathbb{R}^d . The rate of convergence of this difference is described in terms of **Lyapunov exponents**

$$\mathcal{E}_\gamma(\nu, \tilde{\nu}, \omega) = \limsup_{n \rightarrow \infty} \frac{1}{n} \log \|\pi_n^\nu - \pi_n^{\tilde{\nu}}\|,$$

for all $\omega \in \Omega$, where the limit holds in some strong sense (such as probability, mean-square, or almost-surely). However, it turns out that the Lyapunov exponent is almost-surely bounded by a deterministic constant, dependent only on model parameters such as γ .

Before moving on, we define the sequence of stochastic operators $T_n \doteq \psi_n \Lambda^*$, and define the sequence unnormalized posterior distributions,

$$p_n^\nu \doteq T_n p_{n-1}^\nu = T_n T_{n-1} \dots T_1 \nu. \quad (10.4)$$

Clearly, $\pi_n^\nu = p_n^\nu / \langle p_n^\nu, \mathbf{1} \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner-product on \mathbb{R}^d , and $\mathbf{1} = (1, 1, 1, \dots, 1)^*$.

10.2.1 Perron-Frobenius & Oseledec's Theorems

In this section we introduce some general results from matrix theory. The theorems and their usage have heavy dependence on the algebraic concept of an 'exterior product.' For now, we start with a basic theory that we can apply to the eigenvalues and eigenvectors of a Markov chain transition matrix:

Theorem 10.2.1. (Perron-Frobenius). *Assuming that the Markov chain transition matrix Λ is irreducible and of primitive order. Then 1 is a simple eigenvalue of Λ with all other eigenvalues have real-part with absolute value strictly less than one. Moreover, the unique right-eigenvector corresponding to 1 can be multiplied by a constant to equal μ , so that $\lim_n \Lambda^n = (\mu, \mu, \dots, \mu)^*$. Furthermore, for any probability vector ν we have $(\Lambda^*)^n \nu \rightarrow \mu$ as $n \rightarrow \infty$.*

Proof. (see Ethier and Kurtz [23]). □

In filtering, we apply a sequence of matrices T_n which are not time-homogenous. Therefore, we need to consider a more general framework when considering eigenvalues and the space of eigenvectors for T_n . The multiplicative ergodic theorem, also known as Oseledec's theorem, will be useful, but before we present the theorem we need to define the following,

Definition 10.2.1. *An operator $C(x, n)$ where $x = (x_0, \dots, x_n) \in \mathcal{S}^{n+1}$, is a **cocycle** if*

- $C_0(x) = I_{d \times d}$ for all x ,
- $C_n(x) = C_{n-m}(x_m) C_m(x)$.

where $x_m = (x_m, \dots, x_n) \in \mathcal{S}^{n-m+1}$.

Letting $M_n = T_n T_{n-1} \dots T_1$ with the convention that $M_{-1} = I$, we see that M_n is a cocycle. Now, we are ready for Oseledec's theorem:

Theorem 10.2.2. (Oseledec's Multiplicative Ergodic Theorem). *Suppose both M_n and M_n^{-1} are integrable for all n , $\mathbb{E}M_n + \mathbb{E}M_n^{-1} < \infty$ for all $n < \infty$. Then for each $\nu \in \mathbb{R}^d \setminus \{0\}$ we have*

$$\mathcal{V} = \lim_n \frac{1}{n} \log \left(\frac{\|M_n \nu\|}{\|\nu\|} \right) \quad a.s$$

exists and can take up to d -many values. For some $\ell \leq d$, if $\mathcal{V}_1 > \mathcal{V}_2 > \dots > \mathcal{V}_\ell$ are the ℓ -many different (random) limits, then there exist (random) subspaces $\mathbb{R}^d \supsetneq S_\omega^1 \supset S_\omega^2 \supset \dots \supset S_\omega^\ell \supset \{0\}$ such that the limit is \mathcal{V}_i if $\nu \notin S_\omega^i$ for $i \leq \ell$.

Proof. (see page 181 of [12]). □

In Oseledec's theorem, the exponentials $\exp(\mathcal{V}_1) > \exp(\mathcal{V}_2) > \dots > \exp(\mathcal{V}_\ell)$ are the eigenvalues of the matrix $\lim_n (M_n^* M_n)^{1/2n}$. It should be noted that when $\ell < d$, the space spanned by the non-generalized eigenvectors of $\lim_n (M_n^* M_n)^{1/2n}$ will be $\mathbb{R}^\ell \subsetneq \mathbb{R}^d$.

A piece of matrix theory that will be useful is the exterior product. The exterior product or wedge product maps two vectors to the parallelogram that they form. For more than two vectors, the exterior product corresponds to the shape analogous to the parallelogram in the dimension equal to the number of vectors (for instance, the exterior product of three vectors is a parallelepiped). For any two vectors $a, b \in \mathbb{R}^d$, their **exterior product** is

$$a \wedge b = \sum_{i,j} a_i b_j (e_i \wedge e_j)$$

where $(e_i)_i$ represents the canonical basis of \mathbb{R}^d . Some basic properties of the exterior product are

- $e_i \wedge e_j = -e_j \wedge e_i$
- $a \wedge a = 0$
- $\|a \wedge b\|^2 = \|a\|^2 \|b\|^2 - \langle a, b \rangle^2$

and with regard to the Lyapunov exponents from Oseledec's theorem, we have

$$\limsup_n \frac{1}{n} \log \left(\frac{\|M_n(a \wedge b)\|}{\|a \wedge b\|} \right) \leq \mathcal{V}_1 + \mathcal{V}_2 \quad a.s.$$

10.2.2 Lyapunov Exponents of the Filter

Continuing to let $M_n = T_n T_{n-1} \dots T_1$, we can now show that the rate of convergence is given by the spectral gap in M_n :

Theorem 10.2.3. *Assuming that X_n is an ergodic Markov chain, there exists a deterministic function of γ , namely \mathcal{E}_γ , such that for any $\nu \neq \tilde{\nu}$, we have*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \|\pi_n^\nu - \pi_n^{\tilde{\nu}}\| = \mathcal{E}_\gamma(\nu, \tilde{\nu}) \leq \mathcal{E}_\gamma,$$

\mathbb{P} -a.s. In particular, $\mathcal{E}_\gamma = \mathcal{V}_2 - \mathcal{V}_1 < 0$, which is the spectral gap in the matrix $\lim_n ((M_n)^* M_n)^{1/2n}$.

Proof. From the triangle inequality, we can assume W.L.O.G. that the initial distribution of X is its invariant distribution, and take $\tilde{\nu} = \mu$. In this case, the matrices T_n possess a stationary law which is also ergodic. Moreover,

$$\mathbb{E} \log^+ \|T_n\| \leq c \mathbb{E} \max_i \gamma^{-2} (Y_n h(x_i) - .5 h^2(x_i))^+ < \infty.$$

Hence, we can apply Oseledec's theorem to conclude there exists a random subspace S_ω^1 such that if $\nu \notin S_\omega^1$ then

$$\frac{1}{n} \log \|p_n^\nu\| \rightarrow \mathcal{V}_1 \tag{10.5}$$

\mathbb{P} -a.s. In this setting, $\mathcal{V}_1 > \mathcal{V}_2 > \dots > \mathcal{V}_d$ are the Lyapunov exponents associated with the matrix M_n . It is well-known that $((M_n)^* M_n)^{1/2n}$ has a (random) limit a.s., the eigenvalues of which are $e^{\mathcal{V}_i}$. Note that $(M_n)^* M_n$ is a non-negative matrix, thus by Perron-Frobenius theorem the eigenvector associated with the highest eigenvalue of $(M_n)^* M_n$ has all entries real and non-negative. The last property thus holds for $(M_n^* M_n)^{1/2n}$ too, and hence for $\lim_n (M_n^* M_n)^{1/2n}$. Since S_ω^1 must be orthogonal to the eigenvector associated with the highest eigenvalue of $\lim_n (M_n^* M_n)^{1/2n}$, it follows that S_ω^1 cannot include any probability vector with all entries strictly positive. As cases where ν does not have all positive entries, notice that p^ν does for $n \geq k$ where k was the constant such that $\Lambda_{ij}^n > 0$ for $i, j \leq d$ when $n \geq k$. Thus, (10.5) holds for any probability measure $\nu : \mathcal{S} \rightarrow [0, 1]$.

Using Oseledec's theorem again, this time for $\mathbb{R}^d \wedge \mathbb{R}^d$ -valued process $p_n^\mu \wedge p_n^\nu$, there exists a (random) strict subspace $S_\omega^2 \subset \mathbb{R}^d \wedge \mathbb{R}^d$ such that is $\mu \wedge \nu \notin S_\omega^2$ then

$$\frac{1}{n} \log \|p_n^\mu \wedge p_n^\nu\| \rightarrow \mathcal{V}_1 + \mathcal{V}_2 \tag{10.6}$$

\mathbb{P} -a.s., and for $p_n^\mu \wedge p_n^\nu \in S_\omega^2$ we have

$$\limsup_n \frac{1}{n} \log \|p_n^\mu \wedge p_n^\nu\| \leq \mathcal{V}_1 + \mathcal{V}_2 \tag{10.7}$$

\mathbb{P} -a.s.

Then using the inequality $\frac{1}{\sqrt{d}}|\sin(a, b)| \leq \|a - b\| \leq \sqrt{d}|\sin(a, b)|$ where $\sin(a, b)$ is the sine of the angle between vectors a and b (see lemma ??), and the fact that

$$\sin^2(a, b) = 1 - \cos^2(a, b) = \frac{\|a\|^2\|b\|^2 - \langle a, b \rangle^2}{\|a\|^2\|b\|^2} = \frac{\|a \wedge b\|^2}{\|a\|^2\|b\|^2},$$

we can conclude that

$$\begin{aligned} \limsup_n \frac{1}{n} \log \|\pi_n^\mu - \pi_n^\nu\| &= \limsup_n \frac{1}{n} (\log \|p_n^\mu \wedge p_n^\nu\| - \log \|p_n^\mu\| - \log \|p_n^\nu\|) \leq \mathcal{V}_1 + \mathcal{V}_2 - 2\mathcal{V}_1 \\ &= \mathcal{V}_2 - \mathcal{V}_1 < 0, \end{aligned}$$

which completes the proof. \square

The difference $\mathcal{V}_2 - \mathcal{V}_1 < 0$ is a spectral gap and its negativity is sufficient for the stability of the filters. In their paper [4], Atar and Zeitouni proceed to prove that when $\Lambda_{ij} > 0$ for all $i, j \leq d$, then there exists a constant c such that

$$\mathcal{E}_\gamma \leq c < 0$$

where c does not depend on h or γ . They go on to prove that

$$c \leq -2 \min_{i \neq j} \sqrt{\Lambda_{ij}\Lambda_{ji}}.$$

They also prove the following bounds for low-noise models,

$$\limsup_{\gamma \searrow 0} \gamma^2 \mathcal{E}_\gamma \leq -\frac{1}{2} \sum_{i=1}^d \mu_i \min_{i \neq j} (h(x_i) - h(x_j))^2 \quad (10.8)$$

$$\liminf_{\gamma \searrow 0} \gamma^2 \mathcal{E}_\gamma \geq -\frac{1}{2} \sum_{i=1}^d \mu_i \sum_{j=1}^d (h(x_i) - h(x_j))^2. \quad (10.9)$$

Finally, with regard to the ergodic theory, it was shown by Chigansky in 2006 [17] that the Markov-Feller process (X_n, π_n) has a unique invariant measure \mathcal{M} , such that for any continuous g ,

$$\begin{aligned} \lim_n \frac{1}{n} \sum_n g(X_n, \pi_n) &= \sum_i \int g(x_i, u) \mathcal{M}(x_i, du) \\ &= \sum_i \int \mu_i g(x_i, u) \mathcal{M}_{\mu_i}(du) = \lim_n \mathbb{E}g(X_n, \pi_n) \end{aligned} \quad (10.10)$$

where \mathcal{M}_{μ_i} is the μ_i -marginal of \mathcal{M} .

Bibliography

- [1] Ait-Sahalia, Y. (1999) Transition Densities for Interest Rate and Other Derivatives, *The Journal of Finance*, 54(4), pp. 1361–1395.
- [2] A. Alfonsi, “On the Discretization of the CIR (and Bessel squared) Processes,” 2005.
- [3] S. Asmussen, P. Glynn, *Stochastic Simulation: Algorithms and Analysis*. Springer, 2007.
- [4] R. Atar, O. Zeitouni, “Lyapunov Exponents for Finite State Nonlinear Filtering,” SIAM Journal on Control and Optimization, Vol 35, No 1, pp 33-55, January 1997.
- [5] A. Bain, D. Crisan, “Fundamentals of Stochastic Filtering,” Springer 2009.
- [6] Y. Bar-Shalom, X. R. LI, *Estimation and Tracking: Principles, Techniques and Software*. Boston: Artech House, 1993.
- [7] L. Baum, T. Petrie, G. Soules, N. Weiss, “A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains,” The Annals of Mathematical Statistics, Vol. 41, No. 1 (Feb., 1970), pp. 164-171.
- [8] P. Baxendale, P. Chigansky, R. Lipster, “An Asymptotic Theorem for Filtering with Applications to Stability,” Systems & Control Letters Volume 55, Issue 11, November 2006, Pages 908-917.
- [9] T. Björk. *Arbitrage Theory in Continuous Time*. 2nd edition, Oxford Press, 2004.
- [10] B.Z. Bobrovsky, O. Zeitouni “On the joint nonlinear filtering-smoothing of diffusion processes”, Systems & Control Letters Volume 7, Issue 4, July 1986, Pages 317-321
- [11] O. Cappé, E. Moulines, T. Rydén, *Inference in Hidden Markov Models*. Springer 2005.
- [12] R. Carmona, J. Lacroix “Spectral Theory of Random Schrödinger Operators,” Birkhäuser, 1990.
- [13] P. Carr, R. Lee, “Realized Volatility and Variance: Options via Swaps” *Risk*, May 2007.

- [14] P. Carr, D. Madan. “Towards a Theory of Volatility Trading,” reprinted in Option Pricing, Interest Rates, and Risk Management by Musiella, Jouini, Cvitanic, 1998, pp 417–427, University Press.
- [15] P. Carr, L. Wu, “Variance Risk Premia,” March 2005.
- [16] P. Carr, L. Wu, “A Tale of Two Indices,” *The Journal of Derivatives*, Spring 2006.
- [17] P. Chigansky, “An Ergodic Theorem for Filtering Stability with Applications to Stability,” *Systems & Control Letters*, Vol 55, No 11, November 2006, Pages 908-917.
- [18] P. Chigansky, R. Lipster. R. Van Handel, “Intrinsic Methods in Filter Stability,” in *Oxford University Handbook of Nonlinear Filtering* (D. Crisan and B. Rozovsky, eds.), Oxford University Press, to appear.
- [19] E. Derman, K. Demeterfi, M. Kamal, J. Zou, “More Than You Ever Wanted to Know About Volatility Swaps” *Journal of Derivatives* 6(4), 1999, pages 9-32.
- [20] Drăgulescu, A. and Yakovenko, V. (2002) Probability distribution of returns in the Heston model with stochastic volatility, *Quantitative Finance*, 2, pp. 443–453.
- [21] P. Dupuis, H. Kushner *Numerical Methods for Stochastic Control Problems in Continuous Time*, 2nd edition. Springer 2001.
- [22] R.J. Elliot, A.V. Swishchuk. “Pricing Options and Variance Swaps in Markov Modulated Brownian Markets,” *Hidden Markov Models in Finance*, International Series in Operations Research & Management Science, Vol. 104
- [23] Stewart N. Ethier, Thomas G. Kurtz, *Markov Processes: Characterization and Convergence*. Wiley, 1986.
- [24] J.-P. Fouque, G. Papanicolaou, and R. Sircar, Derivatives in Financial Markets with Stochastic Volatility, Cambridge University Press, 2000.
- [25] B.Fristedt, N. Jain, N. Krylov, Filtering and Prediction: A Primer. American Mathematical Society, 2007.
- [26] J. Gatheral, The Volatility Surface, A Practitioner’s Guide. Wiley, 2006.
- [27] S. Howison, A. Rafailidis, H. Rasmussen, “On the pricing and hedging of volatility derivatives” *Applied Mathematical Finance*, Volume 11, Issue 4 December 2000 , pages 317 - 346.
- [28] A. Jazwinski *Stochastic Processes and Filtering Theory*. Dover 1970.

- [29] T. Kailath, A.H. Sayed, B. Hassibi, "Linear Estimation," Prentice Hall, (2000).
- [30] I. Karatzas, S. Shreve, *Brownian Motion and Stochastic Calculus (Graduate Texts in Mathematics)* 2nd Edition, Springer 1992.
- [31] H.Kunita "Asymptotic Behavior of Nonlinear Filtering Errors of Markov Processes," Journal of Multi-Variate Analysis, Vol 1. No 4, pp 365-393, 1971.
- [32] T.G. Kurtz, D.L. Ocone, "Unique Characterization of Conditional Distribution In Nonlinear Filtering," Annals of Probability, 1988, Vol 16, No 1, pp 80-107.
- [33] H. Kushner "Numerical Approximation to Optimal Nonlinear Filters," 2008, <http://www.dam.brown.edu/lcds/publications/>
- [34] Y.A. Kutoyants. Statistical Inference for Ergodic Diffusion Processes. Springer, London, 2004
- [35] S. Mallat, "A Wavelet Tour of Signal Processing," 2nd ed. Academic Press (1999).
- [36] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, Springer 6th edition, 2007.
- [37] E. Pardoux, *Filtrage Non Lineaire et Equations aux Derivees Partielles Stochastiques Associees*. Ecole d'été de Probabilités de Saint-Fleur, 1989.
- [38] E. Platen, N. Bruti-Liberati, *Numerical Solutions of Stochastic Differential Equations in Finance*, Springer 2010.
- [39] Forsyth, David. Ponce, Jean. "Computer Vision – A modern approach"
- [40] L. R. Rabiner "A tutorial on Hidden Markov Models and selected applications in speech recognition". Proceedings of the IEEE 77 (2), (February 1989). pages: 257-286.
- [41] B. Rozovsky, "On Infinite Dimensional Systems of Stochastic Differential Equations Arising in the Theory of Optimal Nonlinear Filtering," Theory of Probability and its Applications, vol 17, no 1, 1972.
- [42] B. Rozovsky, "A simple proof of uniqueness for Kushner and Zakai equations" In Stochastic analysis, ed. E. Mayer-Wolf, 449-458. Boston: Academic Press, 1991.
- [43] B. Rozovsky, A. Petrov, "Optimal Nonlinear Filtering for Track-Before-Detect in IR Image Sequences." SPIE proceedings: Signal and Data Processing of Small Targets, vol. 3809, Denver, Co, 1999.
- [44] B. L. Rozovskii, A. Petrov, R. B. Blazek, "Interacting Banks of Bayesian Matched Filters." SPIE Proceedings: Signal and Data Processing of Small Targets, Vol. 4048, Orlando, FL, 2000.

- [45] Welch, Greg. Bishop, Gary. “An Introduction to the Kalman Filter,” Department of Computer Science, University of North Carolina at Chapel Hill
- [46] R. Whaley, “Derivatives on Market Volatility: Hedging Tools Long Overdue,” Journal of Derivatives, fall 1993.
- [47] G. George Yin, Qing Zhang, *Continuous-Time Markov Chains and Applications: A Singular Perturbation Approach*. Springer, 1998.